arXiv:1510.02451v4 [stat.ME] 15 Jun 2016

The Bouncy Particle Sampler: A Non-Reversible Rejection-Free Markov Chain Monte Carlo Method

Alexandre Bouchard-Côté*, Sebastian J. Vollmer† and Arnaud Doucet†

16th June 2016

*Department of Statistics, University of British Columbia, Canada.

†Department of Statistics, University of Oxford, UK.

Abstract

Many Markov chain Monte Carlo techniques that are available at present rely on discrete-time reversible Markov processes whose transition kernels are variations over the Metropolis–Hastings algorithm. We explore and generalize an alternative scheme recently introduced in the physics literature [23] where the target distribution is explored using a continuous-time non-reversible Markov process. In the Metropolis–Hastings algorithm, a trial move to a region of lower target density, equivalently to higher "energy", than the current state can be rejected with positive probability. In this alternative approach, a particle moves along straight lines around the space and, when facing a high energy barrier, it is not rejected but its path is modified by bouncing against this barrier. By reformulating this algorithm using inhomogeneous Poisson processes, we show that it is possible to exploit standard simulation techniques to simulate exactly this Markov process in a wide range of scenarios of interest to statisticians. Additionally, when the target distribution is given by a product of factors involving only subsets of the state variables, such as the posterior distribution associated with a probabilistic graphical model, this method can be modified to take advantage of this structure by allowing computationally cheaper "local" bounces which only involve the state variables associated to a factor, while the other state variables keep on evolving. In this context, by exploiting the target structure and leveraging techniques from chemical kinetics, we propose several computationally efficient implementations. Experimentally, this new class of Markov chain Monte Carlo schemes compares favorably to state-of-the-art methods on various Bayesian inference tasks, including for high dimensional models and for large data sets.

Keywords: Continuous-time Markov process; Inhomogeneous Poisson process; Markov chain Monte Carlo; Probabilistic graphical models; Rejection-free simulation.

1 Introduction

Markov chain Monte Carlo (MCMC) methods are standard tools to sample from complex high-dimensional probability measures. Many MCMC schemes that are currently available are based on the celebrated Metropolis–Hastings (MH) algorithm and their efficiency is strongly dependent on the ability of the user to design proposal distributions capturing the main features of the target distribution; see [17] for a comprehensive review. We examine, analyze and generalize here a different approach to sample from distributions on $\mathbb{R}^d$ that has been recently proposed in the physics literature [23]. Let the energy be defined as minus the logarithm of an unnormalized case of the target density. In the simplest version of this methodology, a particle is exploring the space

by moving along straight lines and, when it faces a high energy barrier, it bounces against the contour lines of this energy. This non-reversible rejection-free MCMC method will be henceforth referred to as the Bouncy Particle Sampler (BPS). This algorithm has already been adopted to simulate complex physical systems such as hard spheres, polymers and spin models [16, 19, 20, 22]. In those contributions it has been demonstrated experimentally that this scheme can outperform state-of-the-art methods by up to several orders of magnitude.

However, the implementation of the BPS proposed in [23] is not suitable for most target distributions arising from statistics problems as it requires integrating numerically a non-linear function of the gradient of the energy along straight lines, hence introducing approximation errors and perturbing the invariant distribution of the resulting Markov process. In this article we make the following contributions:

Exact simulation schemes: we propose a reformulation of the algorithm proposed by [23] which shows that bounce times correspond to the first arrival time of an inhomogeneous Poisson process. This reformulation allows us to leverage standard simulation techniques [6, Chapter 6] and recent methods from chemical kinetics [27] to obtain new computationally efficient ways to simulate exactly the BPS process for a large class of statistical models.

Factor graphs: when the target distribution can be expressed as a factor graph [28], a representation generalizing graphical models where the target is given by a product of factors and each factor can be a function of only a subset of variables, we show that a physical multi-particle system method discussed in [23, Section III] can be generalized to exploit this representation to achieve additional computational efficiency. This local version of the BPS only manipulates a restricted subset of the components of the states at each bounce but results in a change of all state components, not just the one being updated contrary to standard algorithms such as the Gibbs sampler.

Ergodicity analysis: we present a proof of the ergodicity of BPS when the velocity of the particle is additionally refreshed at the arrival times of a homogeneous Poisson process. When this refreshment step is not carried out, we exhibit a simple counter-example where ergodicity does not hold.

Efficient refreshment: we propose alternative refreshment schemes and compare their computational efficiency experimentally.

We demonstrate empirically that these new MCMC schemes compare favorably to state-of-the-art MCMC methods on various Bayesian inference problems, including in high-dimensional problems and for large data sets. In an extended version of this manuscript, we propose several additional original extensions of the BPS including versions of the algorithm which are applicable to mixed continuous-discrete distributions, distributions restricted to a compact support and a method relying on the use of curved dynamics instead of straight lines. These are not discussed here due to space restrictions.

The rest of this article is organized as follows. In Section 2, we introduce the simplest version of the BPS, propose original ways to implement it and prove that it is ergodic under weak assumptions. Section 3 presents a modification of the original BPS which exploits the factor graph representation of a target distribution and computationally efficient implementations of this scheme. In Section 4, we demonstrate this methodology on various Bayesian models. The proofs are given in the Appendix.

2 The bouncy particle sampler

2.1 Problem statement and notation

Consider a probability distribution π on $\mathbb{R}^d$, equipped with the Borel σ-algebra $\mathcal{B}(\mathbb{R}^d)$. We assume that π admits a probability density with respect to the Lebesgue measure $\mathrm{d}x$ and slightly abuse

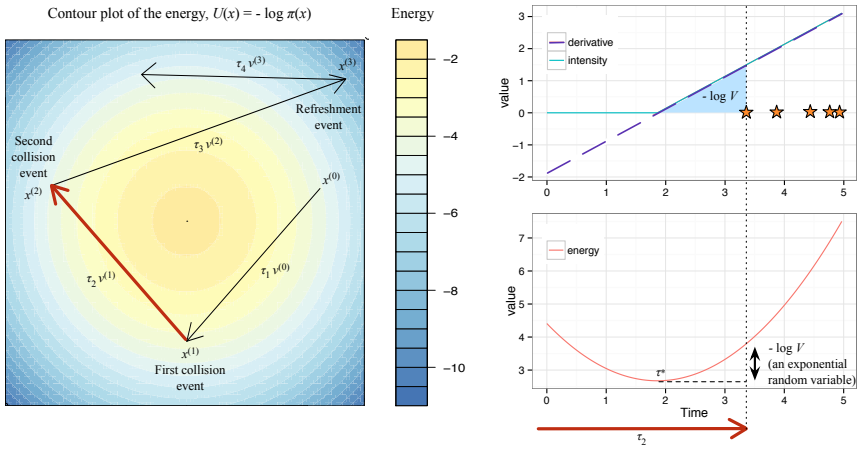

Figure 1: Left: an example of a trajectory produced by the BPS algorithm. The state-space is augmented by a velocity vector v which is maintained constant for a random period of time. After this amount of time, the trajectory can either bounce on the contour line of the energy at the current location (the first two events shown), or undergo velocity refreshment (third event), where the velocity is updated independently of the location. The time at which each bounce event occurs is determined by the first arrival time of an inhomogeneous Poisson process. We show the simulation of the time between the first and second event on the top right. The intensity function is shown in turquoise, and the arrival times are denoted by stars. The intensity function $\chi(t)$ is obtained by computing $\left\langle \nabla U\left(x^{(1)}+v^{(1)}t\right), v^{(1)}\right\rangle$ (purple), and truncating the negative part. On the bottom right, we illustrate a concrete method for computing the arrival times of this Poisson process, described in Example 1.

notation by denoting also this density by π. In most practical scenarios, we only have access to an unnormalized version of this density, that is

$$\pi(x) = \frac{\gamma(x)}{\mathcal{Z}},$$

where $\gamma : \mathbb{R}^d \to (0, \infty)$ can be evaluated point-wise but the normalizing constant $\mathcal{Z} = \int_{\mathbb{R}^d} \gamma(x)\,\mathrm{d}x$ is unknown. We call

$$U(x) = -\log \gamma(x)$$

the associated energy, which is assumed continuously differentiable, and we denote by $\nabla U(x) = \left(\frac{\partial U(x)}{\partial x_1}, \ldots, \frac{\partial U(x)}{\partial x_d}\right)^t$ the gradient of U evaluated at x. We are interested in approximating numerically the expectation of arbitrary test functions $\varphi : \mathbb{R}^d \to \mathbb{R}$ with respect to π.

2.2 Algorithm description

The basic BPS methodology introduced in [23] simulates a continuous trajectory $\{x(t)\}_{t\geq 0}$ in $\mathbb{R}^d$. The trajectory of the particle computed by BPS is composed of piecewise linear segments. Refer to Figure 1 for an example a trajectory generated by BPS on a bivariate isotropic Gaussian target distribution. Each segment in the trajectory is specified by an initial position $x^{(i)} \in \mathbb{R}^d$, a length $\tau_{i+1} \in \mathbb{R}^+$ and a velocity $v^{(i)} \in \mathbb{R}^d$.

Let us denote the times where the velocity changes by $t_i = \sum_{j=1}^{i} \tau_j$ for $i \geq 1$, and set $t_0 = 0$ for convenience. The position at time $t \in [t_i, t_{i+1})$ is thus interpolated linearly, $x(t) = x^{(i)} + v^{(i)}(t - t_i)$, and each segment is connected to the next, $x^{(i+1)} = x^{(i)} + v^{(i)}\tau_{i+1}$. To generate

these segments, the algorithm relies on a position and velocity-dependent intensity function $\lambda : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty)$ and a position-dependent bouncing matrix $R : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ given respectively by

$$\lambda(x, v) = \max\{0, \langle \nabla U(x), v \rangle\} \tag{1}$$

and for any $v \in \mathbb{R}^d$

$$R(x)v = \left(I_d - 2\frac{\nabla U(x)\{\nabla U(x)\}^t}{\|\nabla U(x)\|^2}\right) v = v - 2\frac{\langle \nabla U(x), v \rangle}{\|\nabla U(x)\|^2} \nabla U(x), \tag{2}$$

where I_d denotes the $d \times d$ identity matrix, $\|\cdot\|$ the Euclidean norm, and $\langle w, z \rangle = w^t z$ the scalar product between column vectors w, z.[1] The algorithm also performs a velocity refreshment at random times distributed according to the arrival times of a homogeneous Poisson process of intensity $\lambda^{\text{ref}} \geq 0$, λ^{ref} being a parameter of the BPS algorithm. Throughout the paper, we use the terminology "event" for a time at which either a bounce or a refreshment occurs. The basic version of the BPS algorithm proceeds as follows:

Algorithm 1 Basic BPS algorithm

1. Initialize the state and velocity $\left(x^{(0)}, v^{(0)}\right)$ arbitrarily on $\mathbb{R}^d \times \mathbb{R}^d$.
2. While more events $i = 1, 2, \ldots$ requested do
   (a) Simulate the first arrival time $\tau_{\text{bounce}} \in (0, \infty)$ of an inhomogeneous Poisson process of intensity $\chi(t) = \lambda(x^{(i-1)} + v^{(i-1)}t, v^{(i-1)})$.
   (b) Simulate $\tau_{\text{ref}} \sim \text{Exp}\left(\lambda^{\text{ref}}\right)$.
   (c) Set $\tau_i \leftarrow \min(\tau_{\text{bounce}}, \tau_{\text{ref}})$ and compute the next position
   $$x^{(i)} \leftarrow x^{(i-1)} + v^{(i-1)}\tau_i. \tag{3}$$
   (d) If $\tau_i = \tau_{\text{ref}}$, sample the next velocity $v^{(i)} \sim \mathcal{N}(0_d, I_d)$.
   (e) If $\tau_i = \tau_{\text{bounce}}$, compute the next velocity $v^{(i)}$ using
   $$v^{(i)} \leftarrow R\left(x^{(i)}\right) v^{(i-1)}, \tag{4}$$
   which is the vector obtained once $v^{(i-1)}$ bounces on the plane tangential to the gradient of the energy function at $x^{(i)}$.
3. End While.

In the algorithm above, $\text{Exp}(\delta)$ denotes the exponential distribution of rate δ and $\mathcal{N}(0_d, I_d)$ the standard Gaussian distribution on $\mathbb{R}^d$.[2] Compared to the algorithm described in [23], our formulation of step 2a is expressed in terms of an inhomogeneous Poisson process arrival.

We will show further that the transition kernel of the resulting process $x(t)$ admits π as invariant distribution for any $\lambda^{\text{ref}} \geq 0$ but it can fail to be irreducible when $\lambda^{\text{ref}} = 0$. It is thus critical to use $\lambda^{\text{ref}} > 0$. Our proof of invariance and ergodicity can accommodate some alternative ways to perform the refreshment step 2d. One such variant, which we call restricted refreshment, samples $v^{(i)}$ uniformly on the unit hypersphere $\mathcal{S}^{d-1} = \left\{x \in \mathbb{R}^d : \|x\| = 1\right\}$. We compare experimentally these two variants and others in Section 4.3.

[1] Throughout the paper, when an algorithm contains an expressions of the form $R(x)v$, it is understood that this computation is implemented via the right-hand side of Equation 2 which takes time $O(d)$ rather than the left-hand side, which would naively take time $O(d^2)$.

[2] By exploiting the memorylessness of the exponential distribution, we could alternatively only implement step b) of Algorithm 1 for the i^{th} event when τ_{i-1} corresponds to a bounce and set $\tau_{\text{ref}} \leftarrow \tau_{\text{ref}} - \tau_{i-1}$ otherwise.

2.3 Algorithms for bounce time simulation

Implementing BPS requires sampling the first arrival time τ of a one-dimensional inhomogeneous Poisson process Π of intensity $\chi(t; x, v) = \lambda(x + vt, v)$ given by (1). This intensity depends on the outcome of the last event, but we omit notationally this dependence and write $\chi(t) = \chi(t; x, v)$. This Poisson process Π is denoted by the stars in Figure 1. Simulation from such a process is a well-studied problem; see [6, Chapter 6, Section 1.3]. We briefly review here three standard methods and show how they can be used to implement BPS for examples from Bayesian statistics. The first method described in Section 2.3.1 will be particularly useful when the target is log-concave, while the two others described in Section 2.3.2 and Section 2.3.3 can be applied to more general scenarios.

2.3.1 Simulation using a time-scale transformation

If we let $\Xi(t) = \int_0^t \chi(s)\,\mathrm{d}s$, then the homogeneous Poisson process Π of intensity χ satisfies

$$\mathbb{P}(\tau > u) = \mathbb{P}(\Pi \cap [0, u) = \emptyset) = \exp(-\Xi(u))\frac{(\Xi(u))^0}{0!},$$

and therefore τ can be simulated from a uniform variate $V \sim \mathcal{U}(0, 1)$ via the identity

$$\tau = \Xi^{-1}(-\log(V)), \tag{5}$$

where Ξ^{-1} denotes the quantile function of Ξ, $\Xi^{-1}(p) = \inf\{t : p \leq \Xi(t)\}$. Refer to Figure 1, top right for a graphical illustration. This identity corresponds to the method proposed in [23] to determine the bounce times. Equivalently, this method can be viewed as transforming a homogeneous Poisson process Π_0 of unit rate via an appropriate time-scale transformation [4].

In general, it will not be possible to obtain an analytical expression for Equation (5). However, when the target distribution is strictly log-concave and differentiable, it is easy to use numerical methods to efficiently solve Equation (5) up to numerical precision (see Example 1 below).

Example 1. *Log-concave densities.* Consider the case where the energy is strictly convex (see Figure 1, bottom right). In this case, we can first minimize the energy along the line specified by the initial point and velocity

$$\tau_* = \mathrm{argmin}_{t:t\geq 0}\, U(x + vt),$$

where τ_* is well defined and unique by strict convexity. On the interval $[0, \tau_*)$, which might be empty, we have $\mathrm{d}U(x + vt)/\mathrm{d}t < 0$ and $\mathrm{d}U(x + vt)/\mathrm{d}t \geq 0$ on $[\tau_*, \infty)$. The solution τ of (5) is thus necessarily such that $\tau \geq \tau_*$ and (5) can be rewritten using the gradient theorem as

$$\int_{\tau_*}^{\tau} \frac{\mathrm{d}U(x + vt)}{\mathrm{d}t}\mathrm{d}t = U(x + v\tau) - U(x + v\tau_*) = -\log V.$$

Even if we only compute the energy U point-wise through a black box, we can still efficiently solve this equation through line search within machine precision.

Example 2. *Gaussian distributions.* Consider the scenario where π is a zero-mean multivariate Gaussian distribution of covariance matrix $\frac{1}{2}I_d$ so that $U(x) = \|x\|^2$. Specializing Example 1, simple calculations yield:

$$\tau = \frac{1}{\|v\|^2}\begin{cases} -\langle x, v\rangle + \sqrt{-\|v\|^2 \log V} & \text{if } \langle x, v\rangle \leq 0 \\ -\langle x, v\rangle + \sqrt{\langle x, v\rangle^2 - \|v\|^2 \log V} & \text{otherwise.} \end{cases} \tag{6}$$

2.3.2 Simulation using adaptive thinning

In scenarios where it is difficult to solve (5), the use of a thinning procedure to simulate τ provides another alternative. Assume we have access to local-in-time upper bounds $\bar{\chi}_s(t)$ on $\chi(t)$, that is

$$\bar{\chi}_s(t) = 0 \text{ for all } t < s,$$
$$\bar{\chi}_s(t) \geq \chi(t) \text{ for all } t \geq s$$

and that we can simulate the first arrival time of the inhomogeneous Poisson process $\bar{\Pi}_s$ with intensity $\bar{\chi}_s(t)$ defined on $[s, \infty)$. Algorithm 2 shows the pseudocode for the adaptive thinning procedure.

Algorithm 2 Simulation of the first arrival time of an inhomogeneous Poisson process through thinning

1. Set $s \leftarrow 0$, $\tau \leftarrow 0$.
2. Do
   (a) Set $s \leftarrow \tau$.
   (b) Sample τ as the first arrival point of $\bar{\Pi}_s$ of intensity $\bar{\chi}_s$.
   (c) While $V > \frac{\chi(\tau)}{\bar{\chi}_s(\tau)}$ where $V \sim \mathcal{U}(0,1)$.
3. Return τ.

The event $V > \frac{\chi(\tau)}{\bar{\chi}_s(\tau)}$ corresponds to a rejection step in the thinning algorithm but, in contrast to rejection steps that occur in standard MCMC samplers, in the BPS algorithm this just means that the particle does not bounce and just coasts.

2.3.3 Simulation using superposition

Assume that $U(x)$ can be decomposed as

$$U(x) = \sum_{j=1}^{m} U^{[j]}(x). \tag{7}$$

Under this assumption, if we let $\chi^{[j]}(t) = \max\left(0, \left\langle \nabla U^{[j]}(x + tv), v \right\rangle\right)$ for $j = 1, ..., m$, it follows that

$$\chi(t) \leq \sum_{j=1}^{m} \chi^{[j]}(t).$$

It is therefore possible to use the adaptive thinning algorithm with $\bar{\chi}_s(t) = \bar{\chi}(t) = \sum_{j=1}^{m} \chi^{[j]}(t)$ for $t \geq s$. Moreover, we can simulate from $\bar{\chi}$ via superposition as follows. First, simulate the first arrival time $\tau^{[j]}$ of each inhomogeneous Poisson process with intensity $\chi^{[j]}(t) \geq 0$. Second, return

$$\tau = \min_{j=1,...,m} \tau^{[j]}.$$

Example 3. *Exponential families.* Consider a univariate exponential family with parameter x, observation ι, sufficient statistic $\phi(\iota)$ and log-normalizing constant $A(x)$. If we assume a standard Gaussian prior on x, we obtain the following energy:

$$U(x) = \underbrace{x^2/2}_{U^{[1]}(x)} + \underbrace{-x\phi(\iota)}_{U^{[2]}(x)} + \underbrace{A(x)}_{U^{[3]}(x)},$$

and the bounce intensities are given by:

$$\chi^{[j]}(t) = \max\left(0, \left\langle \nabla U^{[j]}\left(x + vt\right), v\right\rangle\right), j \in \{1,2,3\}.$$

The bounce time $\tau^{[1]}$ for $\chi^{[1]}$ can be computed analytically (see Example 2) whereas the bounce time $\tau^{[2]}$ associated to $\chi^{[2]}$ is given by

$$\tau^{[2]} = \begin{cases} \frac{\log V^{[2]}}{v\phi(\iota)} & \text{if } v\phi(\iota) > 0, \\ +\infty & \text{otherwise,} \end{cases}$$

where $V^{[j]} \sim \mathcal{U}(0,1)$. Finally, the bounce time $\tau^{[3]}$ for $\chi^{[3]}$ is given by

$$\tau^{[3]} = \begin{cases} \tilde{\tau}^{[3]} & \text{if } \tilde{\tau}^{[3]} > 0, \\ +\infty & \text{otherwise,} \end{cases}$$

where $\tilde{\tau}^{[3]} = (A^{-1}(-\log V^{[3]} + A(x)) - x)/v$. For example, with a Poisson distribution with natural parameter x, we obtain

$$\tilde{\tau}^{[3]} = \frac{\log(-\log V^{[3]} + \exp(x)) - x}{v}.$$

Example 4. *Logistic regression.* We consider a binary logistic regression model. We index the data by $r \in \{1, 2, \ldots, R\}$, the class label of the data point r is denoted by c_r and the covariate $k \in \{1, 2, \ldots, d\}$ for data r is denoted $\iota_{r,k}$. We assume for simplicity that $\iota_{r,k} \geq 0$ (this assumption can be easily relaxed but would make the notation more complicated in the following). The parameter vector $x \in \mathbb{R}^d$ is assigned a standard multivariate Gaussian prior distribution, with density denoted by ψ. This yields the following target posterior density:

$$\pi(x) \propto \psi(x) \prod_{r=1}^{R} \frac{\exp(c_r \langle \iota_r, x\rangle)}{1 + \exp\langle \iota_r, x\rangle}. \tag{8}$$

Using the superposition method (Section 2.3.3), simulation of the bounce times can be broken into subproblems corresponding to $R+1$ factors: one factor coming from the prior, with corresponding energy

$$U^{[R+1]}(x) = -\log \psi(x) = \|x\|^2/2 + \text{constant}, \tag{9}$$

and R factors coming from the likelihood of each datapoint, with corresponding energy

$$U^{[r]}(x) = \log(1 + \exp\langle \iota_r, x\rangle) - c_r \langle \iota_r, x\rangle. \tag{10}$$

We denote the corresponding intensities by $\chi^{[r]}$. Simulation of $\chi^{[R+1]}$ is covered in Example 2, so in the following, we focus on the factors $\chi^{[r]}$ for $r \in \{1, 2, \ldots, R\}$. Simulation for those terms can be approached using thinning. In Appendix B.1, we show that :

$$\chi^{[r]}(t) \leq \bar{\chi}^{[r]} = \sum_{k=1}^{d} \mathbf{1}[v_k(-1)^{c_r} \geq 0]\iota_{r,k}|v_k|. \tag{11}$$

Since the bound is constant for a given v, we can sample $\tau^{[r]}$ by simulating an exponential random variable. Sharper bounds could be borrowed from the variational inference literature [15].

2.4 Estimating expectations

Given a realization of the process $x(t)$ over the interval $[0, t_n]$, the expectation $\int_{\mathbb{R}^d} \varphi(x)\,\pi(\mathrm{d}x)$ of a function $\varphi : \mathbb{R}^d \to \mathbb{R}$ with respect to π can be estimated using

$$\frac{1}{t_n}\int_0^{t_n} \varphi(x(t))\,\mathrm{d}t = \frac{1}{t_n}\sum_{i=1}^{n-1}\int_0^{\tau_i} \varphi\left(x^{(i-1)} + v^{(i-1)}s\right)\mathrm{d}s.$$

For example for $\varphi(x) = x_k$, $k \in \{1, 2, \ldots, d\}$, we have

$$\int_0^{\tau_i} \varphi\left(x^{(i-1)} + v^{(i-1)} s\right) \mathrm{d}s = x_k^{(i-1)} \tau_i + v_k^{(i-1)} \frac{\tau_i^2}{2}.$$

When the integral above is intractable, we may just subsample the trajectory of $x(t)$ at regular time intervals to obtain an estimator

$$\frac{1}{L} \sum_{l=0}^{L-1} \varphi(x(l\delta))$$

where $\delta > 0$ and $L = 1 + \lfloor T/\delta \rfloor$. Alternatively, we could approximate these univariate integrals through quadrature.

2.5 Theoretical results

Peters and de With (2012) present an informal proof establishing the fact that the BPS with $\lambda^{\mathrm{ref}} = 0$ admits π as invariant distribution. We provide in Appendix A a rigorous proof of this π-invariance result for $\lambda^{\mathrm{ref}} \geq 0$ and prove that the resulting process is additionally ergodic when $\lambda^{\mathrm{ref}} > 0$. In the following we denote by $P_t(z, \mathrm{d}z')$ the transition kernel of the continuous-time Markov process $z(t) = (x(t), v(t))$.

Proposition 1. *For any $\lambda^{\mathrm{ref}} \geq 0$, the infinitesimal generator associated to the transition kernel P_t of the BPS is given, for any given continuously differentiable functions $h : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$, by*

$$\begin{aligned} \mathcal{L}h(z) &= \lim_{t \to 0} \frac{\int P_t(z, \mathrm{d}z') h(z') - h(z)}{t} \\ &= -\lambda(x, v) h(z) + \langle \nabla_x h, v \rangle + \lambda^{\mathrm{ref}} \int (h(x, v') - h(x, v)) \psi(\mathrm{d}v') \\ &\quad + \lambda(x, v) h(x, R(x) v), \end{aligned}$$

where we recall that $\psi(v)$ denotes the standard multivariate Gaussian density on $\mathbb{R}^d$.

This transition kernel is non-reversible and ρ-invariant, where

$$\rho(z) = \pi(x) \psi(v). \tag{12}$$

If we add the condition $\lambda^{\mathrm{ref}} > 0$, we get the following stronger result.

Theorem 1. *If $\lambda^{\mathrm{ref}} > 0$ then ρ is the unique invariant probability measure of the transition kernel of the BPS and the corresponding process satisfies a strong law of large numbers for ρ-almost every $z(0)$ and $h \in L^1(\rho)$*

$$\lim_{T \to \infty} \frac{1}{T} \int_0^T h(z(t)) \mathrm{d}t = \int h(z) \rho(z) \mathrm{d}z \quad a.s.$$

We exhibit in Section 4.1 a simple example where P_t is not ergodic for $\lambda^{\mathrm{ref}} = 0$.

3 The local bouncy particle sampler

3.1 Structured target distribution and factor graph representation

In numerous applications, the target distribution admits some structural properties that can be exploited by sampling algorithms. For example, the popular Gibbs sampler takes advantages of

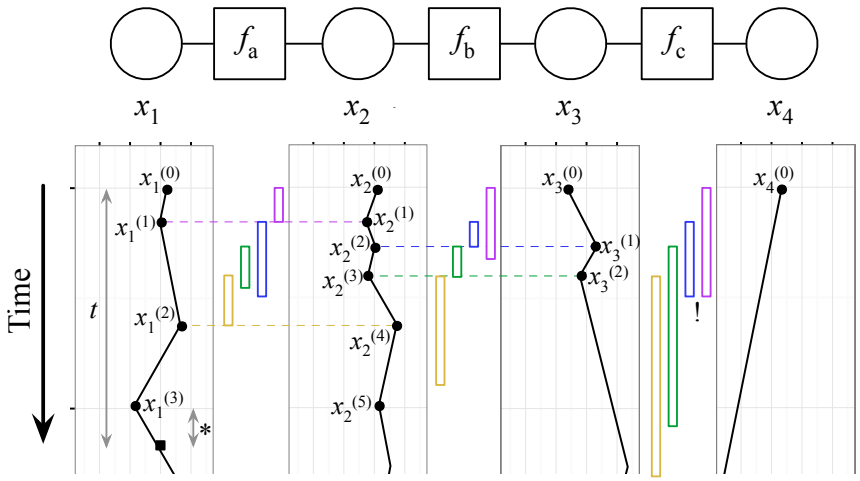

Figure 2: Top: an example of a factor graph with $d = 4$ variables and 3 binary factors, $F = \{f_a, f_b, f_c\}$. Bottom: an example of paths generated by the local BPS algorithm. The circles show the locations that are stored in memory (each with their associated time and velocity after bouncing, not shown). The small black square on the first path is used to demonstrate the execution of Algorithm 3, used to reconstruct the location $x(t)$. The algorithm first identifies $i(t, 1)$, which is 3 in this example. The algorithm then uses the information of the latest event preceding x_t to add to the position at that event, $x_1^{(3)}$, the velocity just after the event, $v_1^{(3)}$, times the time increment denoted by the asterisk, $(t - T_1^{(3)})$. The bottom section of the figure also shows the candidate bounce times used by Algorithm 4 to compute the first four bounce events. The exclamation mark indicates an example where a candidate bounce time need not be recomputed thanks to the sparsity structure of the factor graph.

conditional independence properties. We present here a "local" version of the BPS which can exploit any representation of the target density as a product of positive factors

$$\pi(x) \propto \prod_{f \in F} \gamma_f(x_f) \tag{13}$$

where x_f is a restriction of x to a subset $N_f \subseteq \{1,2,\ldots,d\}$ of the components of x, and F is an index set called the set of factors. Hence the energy associated to π is of the form

$$U(x) = \sum_{f \in F} U_f(x) \tag{14}$$

with $\partial U_f(x)/\partial x_k = 0$ for all variable absent from factor f, i.e. for all $k \in \{1, 2, \ldots, d\} \setminus N_f$.

Such a factorization of the target density can be formalized using factor graphs (Figure 2, top). A factor graph is a bipartite graph, with one set of vertices N called the variables, each corresponding to a component of x ($|N| = d$), and a set of vertices F corresponding to the local factors $(\gamma_f)_{f \in F}$. There is an edge between $k \in N$ and $f \in F$ if and only if $k \in N_f$. Such a representation generalizes undirected graphical models [28, Chap. 2, Section 2.1.3]. For example, factor graphs can have distinct factors connected to the same set of components (i.e. $f \neq f'$ with $N_f = N_{f'}$).

3.2 Local BPS: algorithm description

Similarly to the Gibbs sampler, each step of the local BPS manipulates only a subset of the d components of x. Contrary to the Gibbs sampler, the local BPS does not require sampling from

any conditional distribution and each local calculation results in a change of *all* state components, not just the one being updated—how this can be done implicitly without manipulating the full state at each iteration is described below.

For each factor $f \in F$, we define a local intensity function $\lambda_f : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^+$ and a local bouncing matrix $R_f : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ by

$$\lambda_f(x, v) = \max\{0, \langle \nabla U_f(x), v \rangle\}, \tag{15}$$

$$R_f(x)\, v = v - 2 \frac{\langle \nabla U_f(x), v \rangle \nabla U_f(x)}{\|\nabla U_f(x)\|^2}. \tag{16}$$

It is easy to check that $R_f(x)$ satisfies

$$k \in \{1, 2, \ldots, d\} \setminus N_f \Longrightarrow (R_f(x)v)_k = v_k. \tag{17}$$

When suitable, we will slightly abuse notation and write $R_f(x_f)$ for $R_f(x)$ as $R_f(x_f, x_{-f}) = R_f\left(x_f, x'_{-f}\right)$ for any $x_{-f}, x'_{-f} \in \mathbb{R}^{d-|N_f|}$. Similarly, we will use $\lambda_f(x_f, v_f)$ for $\lambda_f(x, v)$.

We define a collection of Poisson process intensities based on the previous event position $x^{(i-1)}$ and velocity $v^{(i-1)}$: $\chi_f(t) = \chi_f(t; x^{(i-1)}, v^{(i-1)}) = \lambda_f(x^{(i-1)} + v^{(i-1)}t, v^{(i-1)})$. In the local version of BPS, the next bounce time τ is the first arrival of an inhomogeneous Poisson process with intensity $\chi(t) = \sum_{f \in F} \chi_f(t)$. However, instead of modifying all velocity variables at a bounce as in the basic BPS, we first sample a factor f with probability $\chi_f(\tau)/\chi(\tau)$ and modify only the variables connected to the sampled factor in the factor graph. More precisely, the velocity v_f is updated using $R_f(x_f)$ defined in Equation (16). If all the factors are connected to the same variables, the resulting algorithm is an alternative to the method presented in 2.3.3. A generalization of the proof of Proposition 1 given in Appendix A.3 shows that this local BPS algorithm results in a $\pi-$invariant kernel. In the next subsection, we describe various computationally efficient procedures to simulate this process.

For all these implementations, it is useful to encode trajectories in a sparse fashion: each variable $k \in N$ only records information at the times $T_k^{(1)}, T_k^{(2)}, \ldots$ where an event (a bounce or refreshment) affected it. By Equation (17), this represents a sublist of the list of all event times. At each of those times $T_k^{(i)}$, the component's position $x_k^{(i)}$ and velocity $v_k^{(i)}$ right after the event is stored. Let $L_k = L_k^{(0:I)}$ denotes the list of triples $\left((x_k^{(0)}, v_k^{(0)}, 0), (x_k^{(1)}, v_k^{(1)}, T_k^{(1)}), \ldots, (x_k^{(I)}, v_k^{(I)}, T_k^{(I)})\right)$, where $x_k^{(0)}$ and $v_k^{(0)}$ denote the initial position and velocity (see Figure 2, where we show the sparse set of recorded positions). The list $L_k^{(0:I-1)}$ is sufficient to interpolate or extrapolate the position $x_k(t)$. This procedure is detailed in Algorithm 3 and an example is shown in Figure 2. The same procedure can obviously also be used to obtain a subset of node positions, for example those connected to a factor f, denoted $x_f(t)$.

Algorithm 3 Interpolation/extrapolation at time t of the position of node k from a list of events.

1. Find the index $i(t, k)$ associated to the largest time $T_k^{(i)}$ verifying $T_k^{(i)} \leq t$.
2. Return $x_k^{(i(t,k))} + (t - T_k^{(i(t,k))}) v_k^{(i(t))}$.

3.3 Local BPS: efficient implementations

3.3.1 Implementation via priority queue

One way to sample arrivals from a Poisson process with intensity $\chi(t) = \sum_{f \in F} \chi_f(t)$ and selecting a factor f with probability proportional to $\chi_f(t)$ consists of using the superposition method.

We maintain "candidate" bounce times T_f, one for each factor, which are stored in a priority queue Q. This data structure ensures that finding the minimum element can be performed with computational complexity $O(\log |F|)$, and that insertion or update of an element can be performed in $O(\log |F|)$. During the execution of the local algorithm, whereas the lists L_k is used to store past, confirmed events, Q is used to store potential future bounce times (called candidates), only a subset of which will make it to the lists L_k.

We pick the smallest time from the queue Q to determine the next bounce time. When a bounce occurs, a key observation behind efficient implementation of the local BPS is that not all the other candidate bounce times need to be recalculated. Suppose that the bounce was associated with factor f. In this case, only the candidate bounce times $T_{f'}$ corresponding to factors f' with $N_{f'} \cap N_f \neq \emptyset$ need to be recalculated. For example, consider the first bounce in Figure 2 (shown in purple), which is triggered by factor f_a. Then only the velocities for the variables x_1 and x_2 need to be updated. Therefore, only the candidate bounce times for factors f_a and f_b need to be updated. The candidate bounce time for f_c stays constant at that time step (this is shown by an exclamation mark in the figure).

The method is detailed in Algorithm 4.

Algorithm 4 Local BPS algorithm (priority queue implementation)

1. Initialize the state and velocity $\left(x^{(0)}, v^{(0)}\right)$ arbitrarily on $\mathbb{R}^d \times \mathbb{R}^d$.
2. Initialize the global clock $T \leftarrow 0$.
3. For $k \in N$ do
   (a) Initialize the list $L_k \leftarrow \left(x_k^{(0)}, v_k^{(0)}, T\right)$.
4. Set $Q \leftarrow$ **new queue** $\left(x^{(0)}, v^{(0)}, T\right)$.
5. Sample $T_{\text{ref}} \sim \text{Exp}\left(\lambda^{\text{ref}}\right)$.
6. While more events $i = 1, 2, \ldots$ requested do
   (a) $(T, f) \leftarrow$ **select** (**pop**) **smallest candidate bounce time** (Q) (T is the time, f is the associated factor).
   (b) If $T < T_{\text{ref}}$ then
      i. $(v_f)_k \leftarrow v_k^{|L_k|-1}$ for all $k \in N_f$.
      ii. $x_f \leftarrow x_f(T)$ (computed using Algorithm 3).
      iii. For $k \in N_f$ do
         A. $j \leftarrow |L_k|$.
         B. $x_k^{(j)} \leftarrow x_k^{(j-1)} + (T - T_k^{(j-1)}) v_k^{(j-1)}$.
         C. $v_k^{(j)} \leftarrow \{R_f(x_f) v_f\}_k$.
         D. $L_k \leftarrow \left\{L_k, \left(x_k^{(j)}, v_k^{(j)}, T\right)\right\}$ (add the new sample to the list).
      iv. For $f' \in F : N_{f'} \cap N_f \neq \emptyset$ (note: this includes the update of f) do
         A. for all $k \in N_{f'}$.
         B. $x_{f'} \leftarrow x_{f'}(T)$ (computed using Algorithm 3).
         C. Simulate the first arrival time $\tau_{f'}$ of an inhomogeneous Poisson process of intensity $\chi(t) = \lambda_{f'}\left(x_{f'} + t v_{f'}, v_{f'}\right)$ on $[0, +\infty)$.
         D. Set in Q the time associated to f' to the value $T + \tau_{f'}$.
   (c) Else
      i. Sample $v' \sim \mathcal{N}(0_d, I_d)$.
      ii. $Q \leftarrow$ **new queue** (x, v', T_{ref}) where x is the state at time T_{ref} obtained from Algorithm 3.
      iii. Set $T_{\text{ref}} \leftarrow T_{\text{ref}} + \tau_{\text{ref}}$ where $\tau_{\text{ref}} \sim \text{Exp}\left(\lambda^{\text{ref}}\right)$.
7. Return the samples encoded as the lists L_k, $k \in N$.

Algorithm 5 New Queue (x, v, T)

1. For $f \in F$ do
   (a) $(v_f)_k \leftarrow v_k^{|L_k|-1}$ for all $k \in N_f$.
   (b) $x_f \leftarrow x_f(T)$ (computed using Algorithm 3).
   (c) Simulate the first arrival time τ_f of an inhomogeneous Poisson process of intensity $\chi(t) = \lambda_f(x_f + tv_f, v_f)$ on $[0, +\infty)$.
   (d) Set in Q the time associated to f to the value $T + \tau_f$.
2. Return Q

3.3.2 Implementation via thinning

When the number of factors involved in Step 6(b)iv is large, the previous queue-based implementation is potentially computationally expensive. The simplest statistical problem where this arises is when we are interested in sampling the posterior density of a parameter x and have access to a large number R of observations conditionality independent given x. In these cases, we can associate one factor to the prior and one factor for each observation so the total number of factors is $R + 1$ and $x_f = x$ for all $f \in F$.

Implementing BPS in this setup is closely related to the problem of simulating stochastic chemical kinetics and innovative solutions have been proposed in this area. We adapt here the algorithm proposed in [27] to the BPS context. For ease of presentation, we present the algorithm without refreshment and only detail the simulation of the bounce times. This algorithm relies on the ability to compute local-in-time upper bounds on λ_f for all $f \in F$. More precisely, we assume that given a current position x and velocity v, and $\Delta \in (0, \infty]$, we can find a positive number $\bar{\chi}_f$, such that for any $t \in [0, \Delta)$, we have $\bar{\chi}_f \geq \lambda_f(x + vt, v)$. We can also use this method on a subset G of F and combine it with the previously discussed techniques to sample candidate bounce times for factors in $F \backslash G$ but we restrict ourselves to $G = F$ to simplify the presentation.

Algorithm 6 Local BPS algorithm (thinning implementation)

1. Initialize the state and velocity $\left(x^{(0)}, v^{(0)}\right)$ arbitrarily on $\mathbb{R}^d \times \mathbb{R}^d$.
2. Initialize the global clock $T \leftarrow 0$.
3. Initialize $\bar{T} \leftarrow \triangle$ (time until which local upper bounds are valid).
4. Compute local-in-time upper bounds $\bar{\chi}_f$ such that $\bar{\chi}_f \geq \lambda_f(x^{(0)} + v^{(0)}t, v^{(0)})$ for all $t \in [0, \Delta)$ and $f \in F$.
5. While more events $i = 1, 2, \ldots$ requested do
   (a) Sample $\tau \sim \text{Exp}\left(\bar{\chi}\right)$ where $\bar{\chi} = \sum_{f \in F} \bar{\chi}_f$.
   (b) If $\left(T + \tau > \bar{T}\right)$ then
      i. $x^{(i)} \leftarrow x^{(i-1)} + v^{(i-1)}(\bar{T} - T)$.
      ii. $v^{(i)} \leftarrow v^{(i-1)}$.
      iii. For all $f \in F$, update the local-in-time upper bound $\bar{\chi}_f$ to ensure that $\bar{\chi}_f \geq \lambda_f(x^{(i)} + v^{(i)}t, v^{(i)})$ for $t \in [0, \Delta)$.
      iv. Set $T \leftarrow \bar{T}$, $\bar{T} \leftarrow \bar{T} + \triangle$.
   (c) Else
      i. $x^{(i)} \leftarrow x^{(i-1)} + v^{(i-1)}\tau$.
      ii. Sample $\mathcal{F} \in F$ where $\mathbb{P}\left(\mathcal{F} = f\right) = \bar{\chi}_f / \bar{\chi}$. .
      iii. If $V < \lambda_{\mathcal{F}}\left(x^{(i)}, v^{(i-1)}\right) / \bar{\chi}_{\mathcal{F}}$ where $V \sim \mathcal{U}\left(0, 1\right)$ then a bounce for factor $\mathcal{F}$ occurs at time T.
         A. $v^{(i)} \leftarrow R_f\left(x^{(i)}\right) v^{(i-1)}$.
         B. For all $f' \in F : N_{f'} \cap N_{\mathcal{F}} \neq \emptyset$, update the local-in-time upper bound $\bar{\chi}_{f'}$ to ensure that $\bar{\chi}_{f'} \geq \lambda_{f'}(x^{(i)} + v^{(i)}t, v^{(i)})$ for $t \in [0, \bar{T} - T - \tau)$.
      iv. Else
         A. $v^{(i)} \leftarrow v^{(i-1)}$.
      v. Set $T \leftarrow T + \tau$.

Remark 1. This implementation will be particularly useful in scenarios where the operations that involve iterating over F can be done implicitly. More precisely, there are two such operations: summing over the bounds (Step 5a) and sampling over the normalized bounds (Step 5(c)ii). There are situations where it is possible to implement these two operations in constant time as illustrated in Section 4.6 for logistic regression. This method only requires an initial pre-processing of complexity $O\left(|F| \log |F|\right)$ before running the algorithm. Another scenario where sampling quickly from $\mathcal{F}$ is feasible is if the number of distinct upper bounds is much smaller than the number of factors. For example, we only need to sample a factor $\mathcal{F}$ uniformly at random if $b^\star = \bar{\chi}_f = \bar{\chi}_{f'}$ for all f, f' in $F = \{1, \ldots N\}$ and $\bar{\chi} = N \cdot b^\star$ (that is no factor needs to be inspected in order to execute Algorithm 5(c)ii) and the thinning procedure 5c)iii in Algorithm 6 boils down to

$$V \leq \frac{\max\left(0, \left\langle \nabla U_f(x^{(i)}), v^{(i-1)} \right\rangle\right)}{b^\star} = \frac{N \max\left(0, \left\langle \nabla U_f(x^{(i)}), v^{(i-1)} \right\rangle\right)}{\bar{\chi}}. \tag{18}$$

We learned through presentations that a related approach has been adopted in the forthcoming work [3]. In this particular scenario, an alternative local BPS could also be implemented where s factors $\mathcal{F} = (\mathcal{F}_1, \ldots, \mathcal{F}_s)$ are sampled uniformly at random without replacement from $1, ..., N$, the thinning occurs with probability

$$\frac{\frac{N}{s} \max\left(0, \sum_{j=1}^{s} \left\langle \nabla U_{\mathcal{F}_j}(x^{(i)}), v^{(i-1)} \right\rangle\right)}{\bar{\chi}}$$

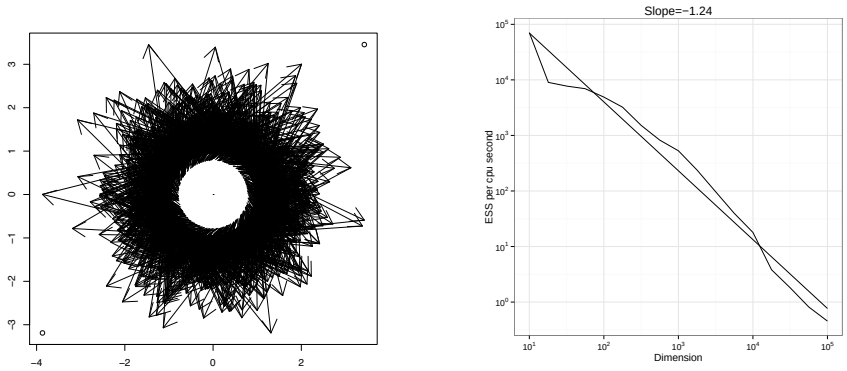

Figure 3: Left: A trajectory of the BPS for $\lambda^{\text{ref}} = 0$, the center of the space is never explored. Right: ESS per CPU second for increasing d for the process with refreshment.

and the bounce is based on $\sum_{j=1}^{s} \left\langle \nabla U_{\mathcal{F}_j}(x^{(i)}), v^{(i-1)} \right\rangle$. It is simple to check that the resulting dynamics preserves π. In contrast to the batchsize $s = 1$, this is not an implementation of local BPS described in Algorithm 6, but instead this corresponds to a local BPS update for a random partition of the factors.

4 Numerical results

4.1 Multivariate Gaussian distributions and the need for refreshment

We use a simple isotropic multivariate Gaussian target distribution, $U(x) = \|x\|^2$, to illustrate the importance of velocity refreshment, restricting ourselves here to BPS with restricted refreshment. Without loss of generality, we assume $\|v^{(i-1)}\| = 1$. From Equation (6), we obtain

$$\left\langle x^{(i)}, v^{(i)} \right\rangle = \begin{cases} -\sqrt{-\log V_i} & \text{if } \left\langle x^{(i-1)}, v^{(i-1)} \right\rangle \leq 0 \\ -\sqrt{\left\langle x^{(i-1)}, v^{(i-1)} \right\rangle^2 - \log V_i} & \text{otherwise} \end{cases}$$

and

$$\left\| x^{(i)} \right\|^2 = \begin{cases} \left\| x^{(i-1)} \right\|^2 - \left\langle x^{(i-1)}, v^{(i-1)} \right\rangle^2 - \log V_i & \text{if } \left\langle x^{(i-1)}, v^{(i-1)} \right\rangle \leq 0 \\ \left\| x^{(i-1)} \right\|^2 - \log V_i & \text{otherwise.} \end{cases}.$$

In particular, these calculations show that if $\left\langle x^{(i)}, v^{(i)} \right\rangle \leq 0$ then $\left\langle x^{(j)}, v^{(j)} \right\rangle \leq 0$ for $j > i$. Using this result, we can show inductively that $\|x^{(i)}\|^2 = \left\|x^{(1)}\right\|^2 - \left\langle x^{(1)}, v^{(1)} \right\rangle^2 - \log V_i$ for $i \geq 2$. In particular for $x^{(0)} = e_1$ and $v^{(0)} = e_2$ with e_i being elements of standard basis of $\mathbb{R}^d$, the norm of the position at all points along the trajectory can never be smaller than 1 as illustrated in Figure 3.

In this scenario, we show that BPS without refreshment admits a countably infinite collection of invariant distributions. Again, without loss of generality, assume $\|v(0)\| = 1$. Let us define $r_t = \|x(t)\|$ and $m_t = \langle x(t), v(t) \rangle / \|x(t)\|$ and denote by χ_k the probability density of the chi distribution with k degrees of freedom.

Proposition 2. *For any dimension $d \geq 2$, the process $(r_t, m_t)_{t \geq 1}$ is Markov and its transition kernel is invariant w.r.t the probability densities $\left\{ f_k(r, m) \propto \chi_k(\sqrt{2}r) \cdot (1 - m^2)^{(k-3)/2}; k \in \{2, 3, 4, \ldots\} \right\}$.*

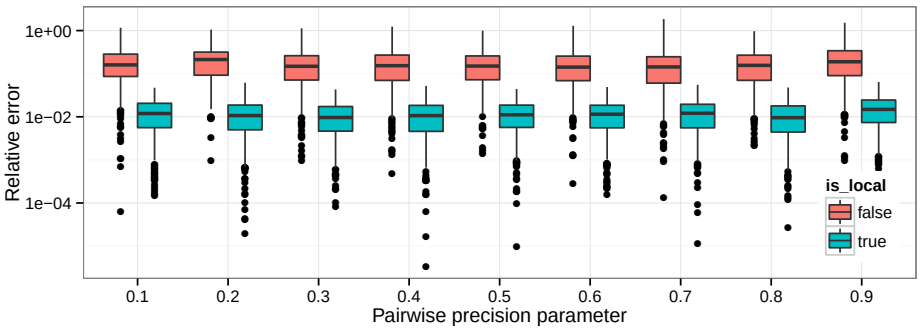

Figure 4: Relative errors for Gaussian chain-shaped random fields. Facets contain results for fields of pairwise precisions 0.1-0.9. Each summarizes the relative errors of 200 (100 local, 100 global) local BPS executions, each ran for a fixed computational budget (a wall-clock time of 60 seconds).

By Theorem 1, we have a unique invariant measure as soon as $\lambda^{\text{ref}} > 0$.

Next, we look at the scaling of the Effective Sample Size (ESS) per CPU second of the basic BPS algorithm as the dimensionality d of the isotropic normal target increases. We use $\lambda_{\text{ref}} = 5$. We focus without loss of generality on the first component of the state vector and estimate the Effective Sample Size (ESS) using the R package *mcmcse* [8] by evaluating the trajectory on a sufficiently fine discretization of the sampled trajectory. The results in log-log scale are displayed in Figure 3. The curve suggests a decay of roughly $d^{-1.24}$, similar to the $d^{-5/4}$ scaling for an optimally tuned Hamiltonian Monte Carlo (HMC) algorithm in a similar same setup [5, Section III], [21, Section 5.4.4]. Both BPS and HMC compare favorably to the d^{-2}scaling of the optimally tuned random walk MH [24].

4.2 Comparison of the global and local schemes

To quantify the potential computational advantage brought by the local version of the algorithm of Section 3 over the global version of Section 2, we compare both algorithms on a sparse Gaussian field. We use a chain-shaped undirected graphical model of length 1000, and perform separate experiments for various pairwise precision parameters for the pairwise interaction between neighbors in the chain. We run the local and global methods for a fixed computational budget (60 seconds), and repeat the experiment in each configuration 100 times. We compute a Monte Carlo estimate of the marginal variance of variable index 500, and compare this estimate to the truth (which can be computed explicitly in this case). The results are shown in Figure 4, in the form of an histogram over relative absolute errors of the 100 executions of each setup. They confirm that the smaller computational complexity per local bounce more than offsets the associate d decrease in expected trajectory segment length. Moreover, the results show that the BPS method is very robust to the pairwise precision used in this sparse Gaussian field model.

4.3 Comparisons of alternative refreshment schemes

In Section 2 and Section 3, the velocity was refreshed using a standard multivariate normal. We compare here this scheme to alternative refreshment schemes:

Global refreshment: sample the entire velocity vector from an standard multivariate Gaussian distribution.

Local refreshment: if the local BPS is being used, the structure specified by the factor graph can be used to design computationally cheaper refreshment operators. We pick one factor $f \in F$ uniformly at random, and we consider resampling only the components of v with

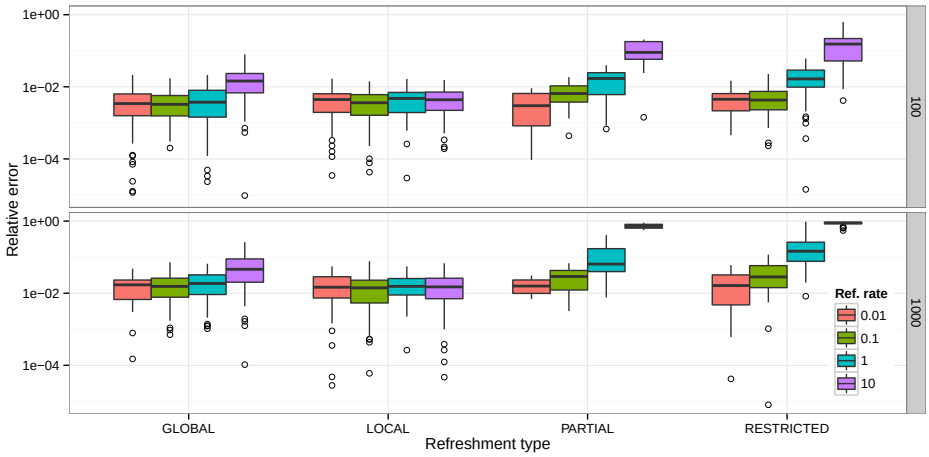

Figure 5: Comparison of four refreshment schemes. The top panel shows results for a 100-dimensional problem, and the bottom one, for a 1000-dimensional problem. The box plots summarize the marginal variance (in log scale) of the variable with index 50 over 100 executions of BPS for each of the refreshment schemes.

indices in N_f. By the same argument used in Section 3, each refreshment will then require bounce time recomputation only for the factors f' with $N_f \cap N_{f'} \neq \emptyset$. Provided that each variable is connected with at least one factor f with $|N_f| > 1$, this scheme is irreducible (and if this condition is not satisfied, additional constant factors can be introduced without changing the target distribution).

Restricted refreshment: this method adds a restriction that the velocities be of unit norm. This scheme corresponds to a different invariant distribution $\rho(x) = \pi(x)\,\phi(v)$ where$\phi(v)$ is now the uniform distribution on $\mathcal{S}^{d-1}$. Refreshment is thus performed by the global refreshment scheme, followed by a re-normalization step.

Restricted partial refreshment: a variant of the restricted refreshment scheme where we sample an angle θ by multiplying a Beta(α, β)-distributed random variable by 2π. We then select a vector uniformly at random from the unit length vectors that have an angle θ from v. We used $\alpha = 1, \beta = 4$ to favor small angles.

The rationale behind the partial refreshment procedure is to suppress the random walk behavior of the particle path arising from a refreshment step independent from the current velocity. Some refreshment is needed to ensure ergodicity but a "good" direction should only be altered slightly. This strategy is akin to the partial momentum refreshment strategy for HMC methods [13], [21, Section 4.3] and, for a normal refreshment, could be similarly implemented. By Lemma 2 all of the above schemes preserve ρ as invariant distribution. We tested these schemes on two versions of the chain-shaped factor graph from the previous section (with the pairwise precision parameter set to 0.5), one with 100 dimensions, and one with 1000 dimensions. All methods are provided with a computational budget of 30 seconds. The results are shown in Figure 5. The results show that the local refreshment scheme is less sensitive to λ^{ref}, performing as well or better than the global refreshment scheme. The performance of the restricted and partial methods appears more sensitive to λ^{ref} and generally inferior to the other two schemes.

4.4 Comparisons with HMC methods on high-dimensional multivariate Gaussian distributions

We compare the local BPS to various state-of-the-art versions of HMC. We use the local refreshment scheme, no partial refreshment and $\lambda^{\text{ref}} = 1$. We select a 100-dimensional Gaussian example from

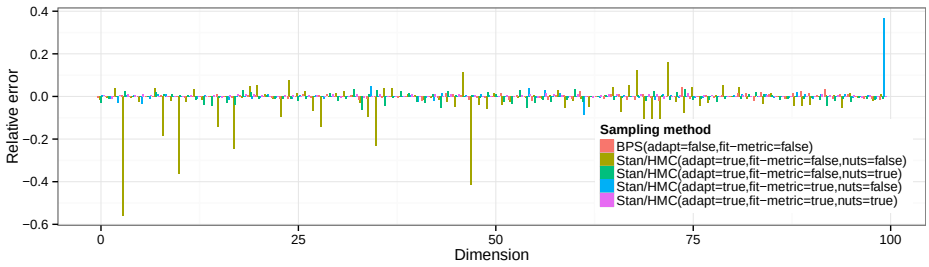

Figure 6: Relative error of marginal variance estimates for a fixed computational budget (30s).

[21, Section 5.3.3.4], where even a basic HMC scheme was shown to perform favorably compared to standard MH methods. We run several methods for this test case, each for a wall clock time of 30 seconds, and measure the relative error on the reconstructed marginal variances. We use Stan [12] as a reference implementation for the HMC algorithms. Different HMC versions are explored by enabling and disabling the NUTS methodology for determining path lengths, and by enabling and disabling adaptive estimation of a diagonal mass matrix. We always exclude the time taken to compile the Stan program in the 30 seconds budget. The three HMC methods tested use 1000 iterations of adaptation, since HMC without adaptation (not shown) yields a zero acceptance rate. In contrast, we use the default value for our local algorithm's tuning parameter ($\lambda^{\text{ref}} = 1$), and no adaptation of the mass matrix. The results (Figure 6) show that this simple implementation performs remarkably well. The adapted HMC performs reasonably well, except for four marginals which are markedly off target. These deviations disappear after incorporating more complex HMC extensions, namely learning a diagonal metric (denoted fit-metric), and adaptively selecting the number of leap-frog steps (denoted nuts).

Next, we perform a series of experiments to investigate the comparative performance of our local method versus NUTS as the dimensionality d increases. Experiments are performed on the chain-shaped Gaussian Random Field of Section 4.2 (with the pairwise precision parameter set to 0.5). We vary the length of the chain (10, 100, 1000), and run Stan's implementation of NUTS for 1000 iterations + 1000 iterations of adaptation. We measure the wall-clock time (excluding the time taken to compile the Stan program) and then run our method for the same wall-clock time. We repeat this 40 times for each chain size. We then measure the absolute value of the relative error on 10 equally spaced marginal variances, and show the rate at which they decrease as the percentage of the samples collected in the fixed computational budget varies from 1 percent to 100 percent. The results are displayed in Figure 7. Note that the gap between the two methods widen as the dimensionality increases. To visualize the different behavior of the two algorithms, we show in Figure 8 three marginals of the Stan and BPS paths from the 100-dimensional example for the first 0.5% of the full trajectories computed in the computational budget.

4.5 Poisson-Gaussian Markov random field

We consider the following hierarchical model. Let $x_{i,j} : i,j \in \{1,2,\dots,10\}$ denote a sparse precision, grid-shaped Gaussian Markov random field with pairwise interactions of the same form as those used in the previous chain examples (pairwise precision set to 0.5). For each i,j, let $y_{i,j}$ be Poisson distributed, independent given $x = (x_{i,j} : i,j \in \{1,2,\dots,10\})$, with rate $\exp(x_{i,j})$. We generate a synthetic dataset from this model and approximate the posterior distribution of x given data $y = (y_{i,j} : i,j \in \{1,2,\dots,10\})$. We run Stan with default settings for $16, 32, 64, \dots, 4096$ iterations. For each configuration, we run local BPS for the same wall-clock time as Stan, using a local refreshment with $\lambda^{\text{ref}} = 1$ and the method from Example 3 to perform the bouncing time computations. We repeat this series of experiments 10 times with different random seeds for the two samplers. We show in Figure 9 estimates of the posterior variances of the variables indexed 0 ($x_{0,0}$), and 50 ($x_{5,0}$). These two marginals are representative of the other variables. Each box plot summarizes the 10 replications. As expected, both methods converge to the same value, but BPS

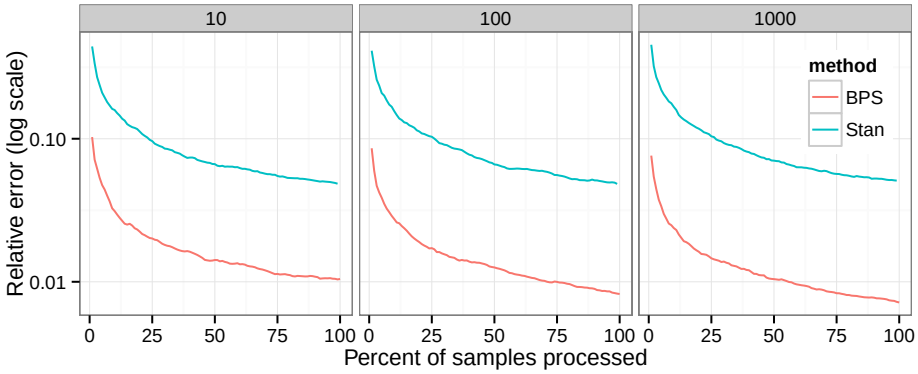

Figure 7: Relative reconstruction error for $d = 10$ (left), $d = 100$ (middle) and $d = 1000$ (right), averaged over 10 of the dimensions and 40 runs. Each panel is ran on a fixed computational budget (corresponding in each panel to the wall clock time taken by 2000 Stan iterations).

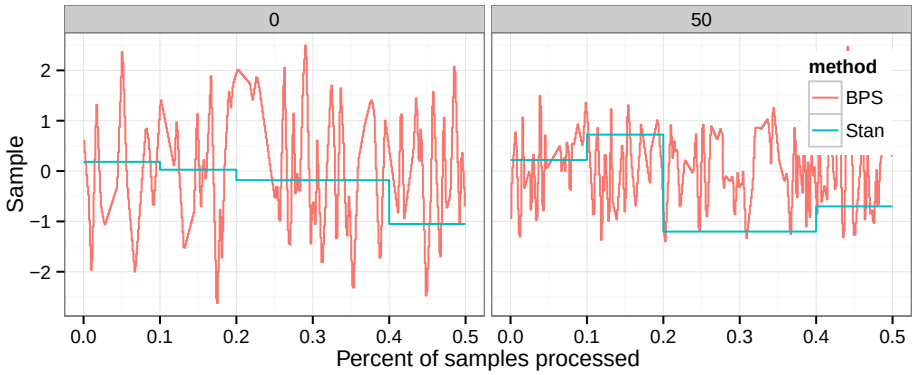

Figure 8: Marginals of the paths for variables index 0 (the left-most variable in the chain) and index 50 (a variable in the middle of the chain). Each of the two marginal paths represent 0.5% of the full trajectory computed in the fixed computational budget used in the 100-dimensional example of Figure 7. While each piecewise constant step in the HMC trajectory is obtained by a sequence of leap-frog steps, the need for a MH step in HMC means that these intermediate steps are not usable Monte Carlo samples. In contrast, the full trajectory obtained from BPS can be used in the calculation of Monte Carlo averages, as explained in Section 2.4.

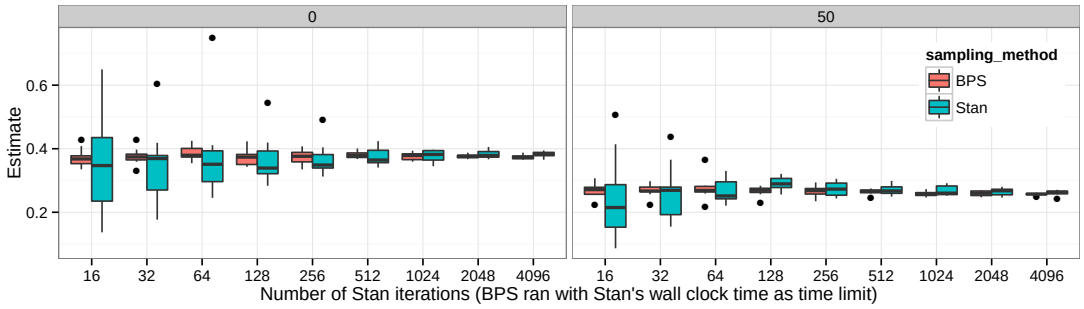

Figure 9: Left: estimates of the posterior variance of $x_{0,0}$; and right, of $x_{5,0}$. Each pair of box-plots is constructed by running Stan with default settings for a set number of iterations, then running BPS with default settings for the same wall-clock time, and repeating this 10 times with different random seeds.

requires markedly less computing time to achieve any given level of accuracy.

4.6 Bayesian logistic regression for large data sets

We consider the logistic regression model introduced in Example 4 in a scenario where a large number of data R is available. In this context, standard MCMC schemes such as the MH algorithm are computationally expensive as they require evaluating the likelihood associated to the R observations at each iteration. This problem has motivated the development of techniques which only evaluate the likelihood of a subset of the data at each iteration. However, most of the methods currently available introduce either some non-vanishing asymptotic bias, e.g. the subsampling MH scheme proposed in [2], or provide consistent estimates converging at a slower rate than regular MCMC algorithms, e.g. the Stochastic Gradient Langevin Dynamics introduced in [30, 26]. To the best of our knowledge, the only available algorithm which only requires evaluating the likelihood of a subset of data at each iteration yet provides consistent estimates converging at the standard Monte Carlo rate is the Firefly algorithm [18].

As alluded to at the beginning of Section 3.3.2, in this context, we can associate $R+1$ factors to the target posterior distribution: one for the prior and one for each data point with $x_f = x$ for all $f \in F$. As outlined in Remark 1 we can provide an alternative to the Firefly algorithm by choosing a data point at random provided a uniform upper bound on the intensity is available (such bound is available for logistic regression see Remark 2 below). In this case Step 5a and Step 5(c)ii can be performed without enumerating the factors.

However, we favor a data dependent rate a strategy instead of a blind strategy that treats all data equally. In the following we outline how this can be achieved in the case of logistic regression provided some precomputation is performed.

After a pre-computation step of complexity $O(R \log R)$ only executed once, we show here that Algorithm 6 can be implemented such that the computational cost of each bounce is independent of R.

We first pre-compute the sum of covariates over the data points, $\iota_k^c = \sum_{r=1}^R \iota_{r,k} \mathbf{1}[c_r = c]$, for each dimensionality $k \in \{1, \dots, d\}$ and class label c. Using these pre-computed quantities, it is then possible to compute the sum $\bar{\chi}$ of the bounding rates $\bar{\chi}^{[r]}$ defined in (11):

$$\bar{\chi} = \sum_{r=1}^{R} \bar{\chi}^{[r]} = \sum_{k=1}^{d} |v_k| \; \iota_k^{\mathbf{1}[v_k < 0]}.$$

If d is large, we can maintain $\bar{\chi}$ in memory and to add and subtract any local updates to it. The implementation of Step 5(c)ii relies on the alias method [6, Section 3.4]. A detailed description of these derivations and of the algorithm is presented in Appendix B.

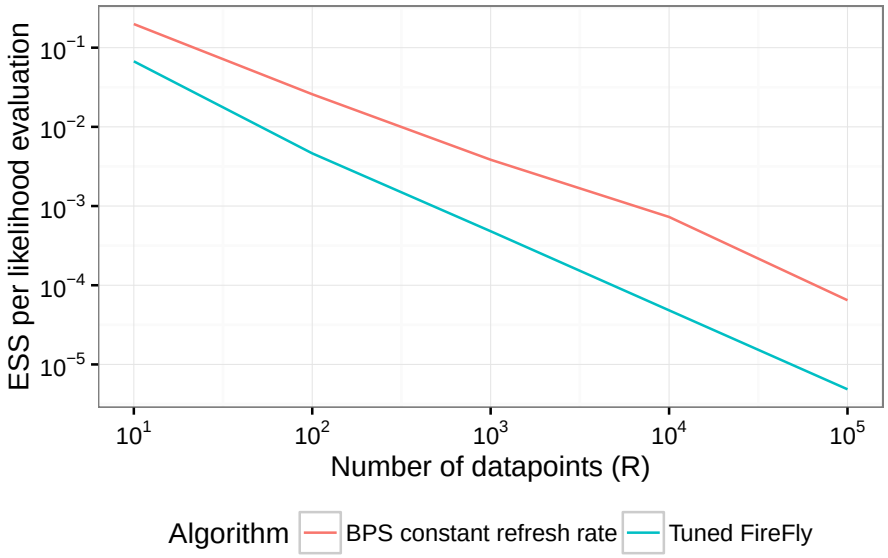

Figure 10: ESS per-datum likelihood evaluation for Local BPS and Firefly.

We compare the local BPS with thinning to the MAP-tuned Firefly algorithm implementation provided by the authors. In [18], it is reported that this version of Firefly outperforms experimentally significantly the standard MH in terms of ESS per-datum likelihood. Local BPS and Firefly are here compared in terms of this criterion, where the ESS is averaged over the $d = 5$ components of x. We generate covariates as $\iota_{rk} \overset{\text{i.i.d.}}{\sim} \mathcal{U}(0.1, 1.1)$ and data $c_r \in \{0, 1\}$ for $r = 1, \ldots, R$ according to (8). We set a standard multivariate normal prior for x. For the algorithm, we set $\lambda_{\text{ref}} = 0.5$ and $\Delta = 0.5$, which is the length of the time interval for which a constant upper bound for the rate associated with the prior is used, see Algorithm 7. Experimentally, local BPS always outperforms Firefly, by about an order of magnitude for large data sets. However, we also observe that both Firefly and local BPS have an ESS per datum likelihood evaluation decreasing in approximately $1/R$ so that the gains brought by these algorithms over a correctly scaled random walk MH algorithm do not appear to increase with R. The rate for local BPS is slightly superior in the regime of up to 10^4 data points, but then returns to the approximate $1/R$ rates. We expect that tighter bounds on the intensities could improve the computational efficiency.

4.7 Bayesian inference of evolutionary parameters

We consider a model from phylogenetics. Given a fixed tree of species with DNA sequences at the leaves, we want to compute the posterior evolutionary parameters encoded into a rate matrix Q. More precisely, we consider an over-parameterized generalized time reversible rate matrix [25] with $d = 10$: 4 unnormalized stationary parameters $x_1, \ldots, x_4$, and 6 unconstrained substitution parameters $x_{\{i,j\}}$, which are indexed by sets of size 2, i.e. where $i, j \in \{1, 2, 3, 4\}$, $i \neq j$. Off-diagonal entries of Q are obtained via $q_{i,j} = \pi_j \exp\left(x_{\{i,j\}}\right)$, where

$$\pi_j = \frac{\exp\left(x_j\right)}{\sum_{k=1}^{4} \exp\left(x_k\right)}.$$

We assign independent standard Gaussian priors on the parameters x_i. We assume that a matrix of aligned nucleotides is provided, where rows are species and columns contains nucleotides believed to come from a shared ancestral nucleotide. Given $x = \left(x_1, \ldots, x_4, x_{\{1,2\}}, \ldots, x_{\{3,4\}}\right)$, and hence Q,

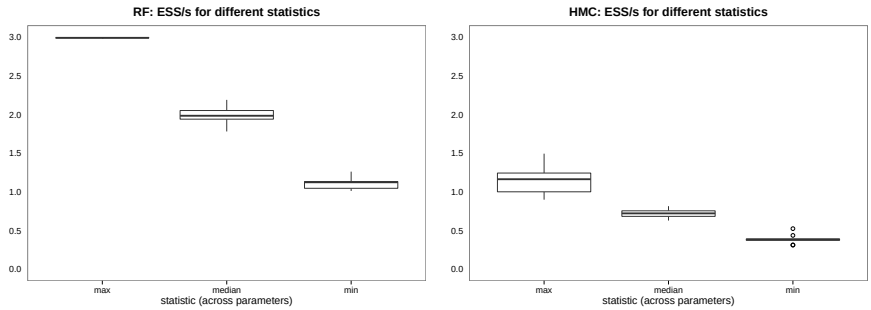

Figure 11: Maximum, median and minimum ESS/s for BPS (left) and HMC (right). The experiments are replicated 10 times with different random seeds.

the likelihood is a product of conditionally independent continuous time Markov chains (CTMC) over {A, C, G, T}, with "time" replaced by a branching process specified by the phylogenetic tree's topology and branch lengths. The parameter x is unidentifiable, and while this can be addressed by bounded or curved parameterizations, the over-parameterization provides an interesting challenge for sampling methods, which need to cope with the strong induced correlations.

We analyze a dataset of primate mitochondrial DNA [11], containing 898 sites and 12 species. We focus on sampling x and fix the tree to a reference tree [14]. We use the basic BPS algorithm with restricted refreshment and $\lambda^{\text{ref}} = 1$ in conjunction with an auxiliary variable-based method similar to the one described in [31], alternating between two moves: (1) sampling CTMC paths along a tree given x using uniformization, (2) sampling x given the path (in which case the derivation of the gradient is simple and efficient). The only difference compared to [31] is that we substitute the HMC kernel by the kernel induced by running BPS. We use this auxiliary variable method because conditioned on the paths, the energy is a convex function and hence we can use the method described in Example 1 to compute the bouncing times.

We compare against a state-of-the-art HMC sampler [29] that uses Bayesian optimization to adapt the key parameters of HMC, the leap-frog stepsize ϵ and trajectory length L, while preserving convergence to the correct target distribution. This sampler was shown to be comparable or better to other state-of-the-art HMC methods such as NUTS. It also has the advantage of having efficient implementations in several languages. We use the author's Java implementation to compare to our Java implementation of the BPS. Both methods view the objective function as a black box (concretely, a Java interface supporting pointwise evaluation and gradient calculation). In all experiments, we initialize at the mode and use a burn-in of 100 iterations and no thinning. The HMC auto-tuner yielded $\epsilon = 0.39$ and $L = 100$. For our method, we use the global sampler and the independent global refreshment scheme.

As a first step, we perform various checks to ensure that both BPS and HMC chains are in close agreement given a sufficiently large number of iterations. We observe that after 20 millions HMC iterations, the highest posterior density intervals from the HMC method are in close agreement with those obtained from BPS (result not shown) and that both method pass the Geweke diagnostic [9].

To compare the effectiveness of the two samplers, we first look at the ESS per second of the model parameters. We show the maximum, median, and maximum over the 10 parameter components, for both BPS and HMC in Figure 11. As observed in Figure 12, the autocorrelation function (ACF) for the BPS decays faster than that of HMC. HMC's slowly decaying ACF is due to the fact that the stepsize ϵ in HMC is selected to be very small by the auto-tuner.

To ensure that the problem does not come from a faulty auto-tuning, we look at the ESS/s for the log-likelihood statistic when varying the stepsize ϵ. The results in Figure 13(right) show that the value selected by the auto-tuner is indeed reasonable, close to the value 0.02 found by brute force maximization. We repeat the experiments with $\epsilon = 0.02$ and obtain the same conclusions. This shows that the problem is genuinely challenging for HMC.

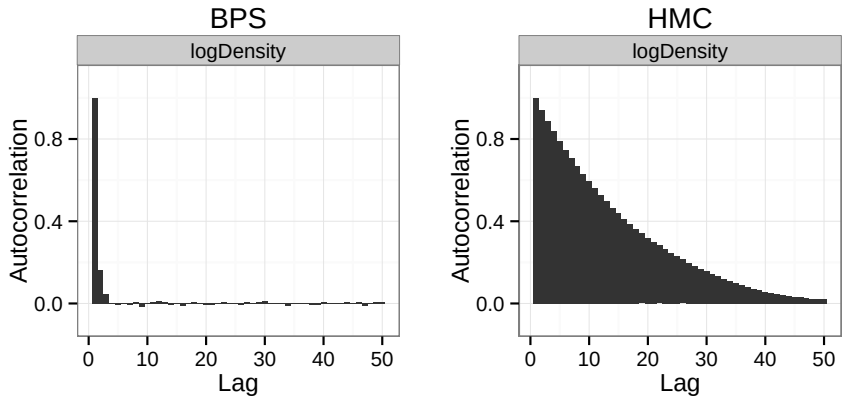

Figure 12: Estimate of the ACF of the log-likelihood statistic for BPS (left) and HMC (right). A similar behavior is observed for the ACF of the other statistics.

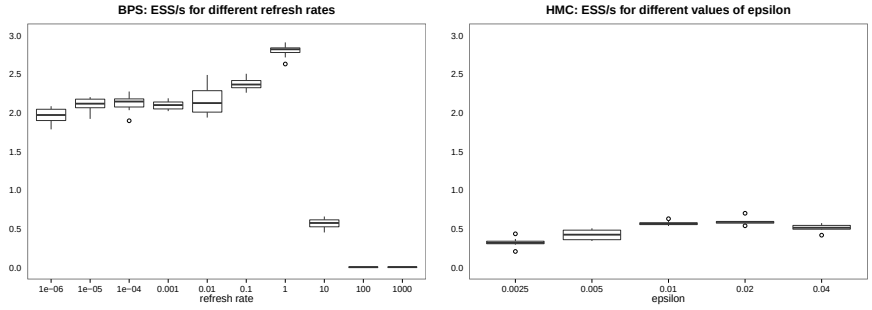

Figure 13: Left: sensitivity of BPS's ESS/s on the log likelihood statistic. Right: sensitivity of HMC's ESS/s on the log likelihood statistic. Each setting is replicated 10 times with different algorithmic random seeds.

The BPS algorithm also exhibits sensitivity to λ^{ref}. We analyze this dependency in Figure 13(left). We observe an asymmetric dependency, where values higher than 1 result in a significant drop in performance, as they bring the sampler closer to random walk behavior. Values one or more orders of magnitudes lower than 1 have a lower detrimental effect. However for a range covering six orders of magnitudes, BPS outperforms HMC at its optimal parameters.

Acknowledgements

Alexandre Bouchard-Côté's research is partially supported by a Discovery Grant from the National Science and Engineering Research Council. Arnaud Doucet's research is partially supported by the Engineering and Physical Sciences Research Council, grants EP/K000276/1, EP/K009850/1. Sebastian Vollmer's research is partially supported by the grants EP/K009850/1 and EP/N000188/1. We thank Markus Upmeier for helpful discussions of differential geometry as well as Nicholas Galbraith and Tingting Zhao for their comments.

References

[1] R.J. Adler and J.E. Taylor. *Topological Complexity of Smooth Random Functions*, volume 2019 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 2011. École d'Été de Probabilités de Saint-Flour XXXIX.

[2] R. Bardenet, A. Doucet, and C. Holmes. Towards scaling up MCMC: An adaptive subsampling approach. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

[3] J. Bierkens, P. Fernhead, and G. O. Roberts. Zig zag monte carlo. in preparation, 2016.

[4] P. Bratley, B.L. Fox, and L.E. Schrage. *A Guide to Simulation*. Springer-Verlag, Berlin, 1983.

[5] M. Creutz. Global Monte Carlo algorithms for many-fermion systems. *Physical Review D*, 38(4):1228–1237, 1988.

[6] L. Devroye. *Non-uniform Random Variate Generation*. Springer-Verlag, New York, 1986.

[7] M. Einsiedler and T. Ward. *Ergodic Theory: with a view towards Number Theory*, volume 259 of *Graduate Texts in Mathematics*. Springer-Verlag, London, 2011.

[8] J.M. Flegal, J. Hughes, and D. Vats. *mcmcse: Monte Carlo Standard Errors for MCMC*, 17-08-2015. R package version 1.1-2.

[9] J. Geweke. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. *Bayesian Statistics*, 4:169–193, 1992.

[10] M. Hairer. Convergence of Markov processes. http://www.hairer.org/notes/Convergence.pdf, 2010. Lecture notes, University of Warwick.

[11] K. Hayasaka, Takashi G., and Satoshi H. Molecular phylogeny and evolution of primate mitochondrial DNA. *Molecular Biology and Evolution*, 5:626–644, 1988.

[12] M.D. Hoffman and A. Gelman. The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(4):1593–1623, 2014.

[13] A.M. Horowitz. A generalized guided Monte Carlo algorithm. *Physics Letters B*, 268(2):247–252, 1991.

[14] J.P. Huelsenbeck and F. Ronquist. MRBAYES: Bayesian inference of phylogenetic trees. *Bioinformatics*, 17(8):754–755, 2001.

[15] T. S. Jaakkola and M. I. Jordan. Bayesian logistic regression: a variational approach. *Statistics and Computing*, 10:25–37, 2000.

[16] T.A. Kampmann, H.H. Boltz, and J. Kierfeld. Monte Carlo simulation of dense polymer melts using event chain algorithms. *Journal of Chemical Physics*, (143):044105, 2015.

[17] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, 2008.

[18] D. Maclaurin and R.P. Adams. Firefly Monte Carlo: Exact MCMC with subsets of data. In *Uncertainty in Artificial Intelligence*, volume 30, pages 543–552, 2014.

[19] M. Michel, S.C. Kapfer, and W. Krauth. Generalized event-chain Monte Carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps. *Journal of Chemical Physics*, 140(5):054116, 2014.

[20] M Michel, J. Mayer, and W. Krauth. Event-chain Monte Carlo for classical continuous spin models. *Europhysics Letters*, (112):20003, 2015.

[21] R.M. Neal. Markov chain Monte Carlo using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, 2011.

[22] Y. Nishikawa, M. Michel, W. Krauth, and K. Hukushima. Event-chain Monte Carlo algorithm for the Heisenberg model. *Physical Review E*, (92):063306, 2015.

[23] E.A. J. F. Peters and G. de With. Rejection-free Monte Carlo sampling for general potentials. *Physical Review E*, 85:026703, 2012.

[24] G. O. Roberts and J.S. Rosenthal. Optimal Scaling for Various Metropolis-Hastings Algorithms. *Statistical Science*, 16(4):351–367, 2001.

[25] S. Tavaré. Some probabilistic and statistical problems in the analysis of DNA sequences. *Lectures on Mathematics in the Life Sciences*, 17:56–86, 1986.

[26] Y. W. Teh, A. H. Thiéry, and S. J. Vollmer. Consistency and fluctuations for stochastic gradient Langevin dynamics. *Journal of Machine Learning Research*, 2015.

[27] V.H. Thanh and C. Priami. Simulation of biochemical reactions with time-dependent rates by the rejection-based algorithm. *Journal of Chemical Physics*, 143(5):054104, 2015.

[28] M.J. Wainwright and M.I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.

[29] Z. Wang, S. Mohamed, and N. De Freitas. Adaptive Hamiltonian and Riemann manifold Monte Carlo. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1462–1470, 2013.

[30] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning*, pages 681–688, 2011.

[31] T Zhao, Z. Wang, A. Cumberworth, J. Gsponer, N. De Freitas, and A. Bouchard-Côté. Bayesian analysis of continuous time Markov chains with application to phylogenetic modelling. *Bayesian Analysis*, 2015.

A Proofs of Section 2

A.1 Proof of invariance w.r.t ρ

Let μ_t be the law of $z(t)$. In the following, we prove invariance by explicitly verifying that the time evolution of the density $\frac{\mathrm{d}\mu_t}{\mathrm{d}t} = 0$ is zero if the initial distribution μ_0 is given by (12). This

is achieved by deriving the forward Kolmogorov equation describing the evolution of the marginal density of the stochastic process. For simplicity, we start by presenting the invariance argument when $\lambda^{\text{ref}} = 0$.

Notation and description of the algorithm. We denote a pair of position and velocity by $z = (x, v) \in \mathbb{R}^d \times \mathbb{R}^d$ and we denote translations by $\Phi_t(z) = (\Phi_t^{\text{pos}}(z), \Phi_t^{\text{dir}}(z)) = (x + vt, v)$. The time of the first bounce coincides with the first arrival T_1 of an inhomogeneous Poisson process with intensity $\chi(t) = \lambda(\Phi_t(z))$ where:

$$\lambda(z) = \max\left\{0, \left\langle \nabla U\left(x\right), v\right\rangle\right\}. \tag{19}$$

It follows that the probability of having no bounce in the interval $[0, t]$ is given by:

$$\text{No}_t(z) = \exp\left(-\int_0^t \lambda(\Phi_s(z))\mathrm{d}s\right), \tag{20}$$

and the density of the random variable T_1 is given by:

$$\begin{aligned} q(t_1; z) = &\quad \mathbf{1}[t_1 > 0]\tfrac{\mathrm{d}}{\mathrm{d}t_1}\left(1 - \text{No}_{t_1}(z)\right) &(21)\\ = &\quad \mathbf{1}[t_1 > 0]\text{No}_{t_1}(z)\lambda(\Phi_{t_1}(z)). &(22)\end{aligned}$$

If a bounce occurs, then the algorithm follows a translation path for time T_1, at which point the velocity is updated using a bounce operation $C(z)$, defined as:

$$C\left(z\right) = \left(x, R\left(x\right)v\right) \tag{23}$$

where

$$R\left(x\right)v = v - 2\frac{\left\langle \nabla U\left(x\right), v\right\rangle \nabla U\left(x\right)}{\left\|\nabla U\left(x\right)\right\|^2}. \tag{24}$$

The algorithm then continues recursively for time $t - T_1$, in the following sense: a second bounce time T_2 is simulated by adding to T_1 a random increment with density $q(\cdot; C \circ \Phi_{t_1}(z))$. If $T_2 > t$, then the output of the algorithm is $\Phi_{t-t_1} \circ C \circ \Phi_{t_1}(z)$, otherwise an additional bounce is simulated, etc. More generally, given an initial point z and a sequence $\mathbf{t} = (t_1, t_2, \dots)$ of bounce times, the output of the algorithm at time t is given by:

$$\Psi_{\mathbf{t},t}\left(z\right) = \begin{cases} \Phi_t(z) & \text{if } t_1 > 0 \text{ or } \mathbf{t} = (\,), \\ \Psi_{\mathbf{t}',t-t_1}\left(z\right) \circ C \circ \Phi_{t_1}(z) & \text{otherwise,} \end{cases} \tag{25}$$

where $(\,)$ denotes the empty list and $\mathbf{t}'$ the suffix of $\mathbf{t}$: $\mathbf{t}' = (t_2, t_3, \dots)$. As for the bounce times, they are distributed as follows:

$$\begin{aligned} T_1 &\sim q(\,\cdot\,; z) &(26)\\ T_i - T_{i-1} | T_{1:i-1} &\sim q\Big(\,\cdot\,;\ \underbrace{\Psi_{T_{1:i-1},T_{i-1}}(z)}_{\text{Pos. after collision } i-1}\ \Big), \quad i \in \{2, 3, 4, \dots\} &(27)\end{aligned}$$

where $T_{1:i-1} = (T_1, T_2, \dots, T_{i-1})$.

Decomposition by the number of bounces. Let h denote an arbitrary non-negative measurable test function. We show how to decompose expectations of the form $\mathbb{E}[h(\Psi_{\mathbf{T},t}(z))]$ by the number of bounces in the interval $(0, t)$. To do so, we introduce a function $\#\text{Col}_t(\mathbf{t})$, which returns the number of bounces in the interval $(0, t)$:

$$\#\text{Col}_t(\mathbf{t}) = \min\left\{n \geq 1 : t_n > t\right\} - 1. \tag{28}$$

From this, we get the following decomposition:

$$\begin{aligned} \mathbb{E}[h(\Psi_{\mathbf{T},t}(z))] &= \mathbb{E}[h(\Psi_{\mathbf{T},t}(z)) \sum_{n=0}^{\infty} \mathbf{1}[\#\text{Col}_t(\mathbf{T}) = n]] &(29)\\ &= \sum_{n=0}^{\infty} \mathbb{E}[h(\Psi_{\mathbf{T},t}(z))\mathbf{1}[\#\text{Col}_t(\mathbf{T}) = n]]. &(30)\end{aligned}$$

On the event that no bounce occurs in the interval $[0,t)$, i.e. $\#\mathrm{Col}_t(\mathbf{T})=0$, the function $\Psi_{\mathbf{T},t}(z)$ is equal to $\Phi_t(z)$, therefore:

$$\begin{aligned}\mathbb{E}[h(\Psi_{\mathbf{T},t}(z))\mathbf{1}[\#\mathrm{Col}_t(\mathbf{T})=0]] &= h(\Phi_t(z))\mathbb{P}(\#\mathrm{Col}_t(\mathbf{T})=0) &(31)\\ &= h(\Phi_t(z))\mathrm{No}_t(z). &(32)\end{aligned}$$

Indeed, on the event that $n\geq 1$ bounces occur, the random variable $h(\Phi_t(z))$ only depends on a finite dimensional random vector, $(T_1,T_2,\dots,T_n)$, so we can write the expectation as an integral with respect to the density $\widetilde{q}(t_{1:n};t,z)$ of these variables:

$$\begin{aligned}&\mathbb{E}[h(\Psi_{\mathbf{T},t}(z))\mathbf{1}[\#\mathrm{Col}_t(\mathbf{T})=n]] &(33)\\ &=\mathbb{E}\left[h(\Psi_{\mathbf{T},t}(z))\mathbf{1}[0<T_1<\cdots<T_n<t<T_{n+1}]\right]\\ &=\int\cdots\int_{0<t_1<\cdots<t_n<t<t_{n+1}} h(\Psi_{t_{1:n},t}(z))q(t_1;z)\prod_{i=2}^{n+1}q(t-t_{i-1};\Psi_{t_{1:i-1},t_{i-1}}(z))\mathrm{d}t_{1:n+1}\\ &=\int\cdots\int_{0<t_1<\cdots<t_n<t} h(\Psi_{t_{1:n},t}(z))\widetilde{q}(t_{1:n};t,z)\mathrm{d}t_{1:n}, &(34)\end{aligned}$$

where:

$$\widetilde{q}(t_{1:n};t,z)=q(t_1;z)\times\begin{cases}\mathrm{No}_{t-t_1}(\Phi_{t_1}(z)) & \text{if } n=1\\ \mathrm{No}_{t-t_n}(\Phi_{t_{1:n},t_n}(z)) & \prod_{i=2}^n q\left(t_i-t_{i-1};\Psi_{t_{1:i-1},t_{i-1}}(z)\right) \text{ if } n\geq 2.\end{cases}$$

To include Equations (32) and (34) under the same notation, we define $t_{1:0}$ to the empty list, $(\,)$, $\widetilde{q}((\,);t,z)=\mathrm{No}_t(z)$, and abuse the integral notation so that for all $n\in\{0,1,2,\dots\}$:

$$\mathbb{E}[h(\Psi_{\mathbf{T},t}(z))\mathbf{1}[\#\mathrm{Col}_t(\mathbf{T})=n]]=\int\cdots\int_{0<t_1<\cdots<t_n<t} h(\Psi_{t_{1:n},t}(z))\widetilde{q}(t_{1:n};t,z)\mathrm{d}t_{1:n}. \quad (35)$$

Marginal density. Let us fix some arbitrary time $t>0$. We seek a convenient expression for the marginal density at time t, $\mu_t(z)$, given an initial vector $Z\sim\rho$, where ρ is the hypothesized stationary density (12) on $\mathcal{Z}$. To do so, we look at the expectation of an arbitrary non-negative measurable test function h:

$$\begin{aligned}\mathbb{E}[h(\Psi_{\mathbf{T},t}(Z))] &= \mathbb{E}\Big[\mathbb{E}[h(\Psi_{\mathbf{T},t}(Z))|Z]\Big] &(36)\\ &= \sum_{n=0}^{\infty}\mathbb{E}\Big[\mathbb{E}[h(\Psi_{\mathbf{T},t}(Z))\mathbf{1}[\#\mathrm{Col}_t(\mathbf{T})=n]|Z]\Big] &(37)\\ &= \sum_{n=0}^{\infty}\int_{\mathcal{Z}}\rho(z)\int\cdots\int_{0<t_1<\cdots<t_n<t}h(\Psi_{t_{1:n},t}(z))\widetilde{q}(t_{1:n};t,z)\mathrm{d}t_{1:n}\mathrm{d}z &(38)\\ &= \sum_{n=0}^{\infty}\int\cdots\int_{0<t_1<\cdots<t_n<t}\int_{\mathcal{Z}}\rho(z)h(\Psi_{t_{1:n},t}(z))\widetilde{q}(t_{1:n};t,z)\mathrm{d}z\mathrm{d}t_{1:n} &(39)\\ &= \sum_{n=0}^{\infty}\int\cdots\int_{0<t_1<\cdots<t_n<t}\int_{\mathcal{Z}}\rho(\Psi^{-1}_{t_{1:n},t}(z'))h(z')\widetilde{q}(t_{1:n};t,\Psi^{-1}_{t_{1:n},t}(z'))\left|\det D\Psi^{-1}_{t_{1:n},t}\right|\mathrm{d}z'\mathrm{d}t_{1:n}\\ &= \int_{\mathcal{Z}}h(z')\underbrace{\sum_{n=0}^{\infty}\int\cdots\int_{0<t_1<\cdots<t_n<t}\rho(\Psi^{-1}_{t_{1:n},t}(z'))\widetilde{q}(t_{1:n};t,\Psi^{-1}_{t_{1:n},t}(z'))\mathrm{d}t_{1:n}}_{\mu_t(z')}\,\mathrm{d}z'. &(40)\end{aligned}$$

We used the following in the above derivation successively the law of total expectation, equation (30), equation (36), Tonelli's theorem and the change of variables, $z'=\Psi_{t_{1:n},t}(z)$, justified since for any fixed $0<t_1<t_2<\cdots<t_n<t<t_{n+1}$, $\Psi_{t_{1:n},t}(\cdot)$ is a bijection (being a composition of bijections). Now the absolute value of the determinant is one since $\Psi_{\mathbf{t},t}(z)$ is a composition of unit-Jacobian mappings and, by using Tonelli's theorem again, we obtain that the expression above the brace is necessarily equal to $\mu_t(z')$ since h is arbitrary.

Derivative. Our goal is to show that for all $z'\in\mathcal{Z}$

$$\frac{\mathrm{d}\mu_t(z')}{\mathrm{d}t}=0.$$

Since the process is time homogeneous, once we have computed the derivative, it is enough to show that it is equal to zero at $t = 0$. To do so, we decompose the computation according to the terms I_n in Equation (40):

$$\mu_t(z') = \sum_{n=0}^{\infty} I_n(z', t) \tag{41}$$

$$I_n(z', t) = \int \cdots \int_{0<t_1<\cdots<t_n<t} \rho(\Psi^{-1}_{t_{1:n},t}(z'))\widetilde{q}(t_{1:n}; t, \Psi^{-1}_{t_{1:n},t}(z'))\mathrm{d}t_{1:n}. \tag{42}$$

The categories of terms in Equation (41) to consider are:

No bounce: $n = 0$, $\Psi_{t_{1:n},t}(z) = \Phi_t(z)$, or,

Exactly one bounce: $n = 1$, $\Psi_{t_{1:n},t}(z) = F_{t,t_1} := \Phi_{t-t_1} \circ C \circ \Phi_{t_1}(z)$ for some $t_1 \in (0, t)$, or,

Two or more bounces: $n \geq 2$, $\Psi_{t_{1:n},t}(z) = \Psi_{t-t_2} \circ C \circ F_{t_2,t_1}(z)$ for some $0 < t_1 < t_2 < t$

In the following, we show that the derivative of the terms in the third category, $n \geq 2$, are all equal to zero, while the derivative of the first two categories cancel each other.

No bounce in the interval. From Equation (32):

$$I_0(z', t) = \rho(\Phi_{-t}(z'))\mathrm{No}_t(\Phi_{-t}(z')). \tag{43}$$

We now compute the derivative at zero of the above expression:

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t} I_0(z', t)\Big|_{t=0} = \mathrm{No}_0 \quad & (\Phi_0(z')) \left.\frac{\mathrm{d}\rho(\Phi_{-t}(z'))}{\mathrm{d}t}\right|_{t=0} + \\ & \rho(\Phi_0(z')) \left.\frac{\mathrm{dNo}_t(\Phi_{-t}(z'))}{\mathrm{d}t}\right|_{t=0}\end{aligned} \tag{44}$$

The first term in the above equation can be simplified as follows:

$$\mathrm{No}_0(\Phi_0(z'))\frac{\mathrm{d}\rho(\Phi_{-t}(z'))}{\mathrm{d}t} = \frac{\mathrm{d}\rho(\Phi_{-t}(z'))}{\mathrm{d}t} \tag{45}$$

$$= \left\langle \frac{\partial\rho(\Phi_{-t}(z'))}{\partial\Phi^{\mathrm{pos}}_{-t}(z')}, \frac{\mathrm{d}\Phi^{\mathrm{pos}}_{-t}(z')}{\mathrm{d}t} \right\rangle + \left\langle \frac{\partial\rho(\Phi_{-t}(z'))}{\partial\Phi^{\mathrm{dir}}_{-t}(z')}, \underbrace{\frac{\mathrm{d}\Phi^{\mathrm{dir}}_{-t}(z')}{\mathrm{d}t}}_{=\mathbf{0}} \right\rangle \tag{46}$$

$$= \left\langle \frac{\partial\rho(z)}{\partial x}, -v' \right\rangle \tag{47}$$

$$= \left\langle \frac{\partial}{\partial x}\frac{1}{Z}\exp\left(-U(x)\right)\psi\left(v\right), -v' \right\rangle$$

$$= \rho(\Phi_{-t}(z'))\left\langle \nabla U(x), v' \right\rangle, \tag{48}$$

where $x = \Phi^{\mathrm{pos}}_{-t}(z')$. The second term in Equation (44) is equal to:

$$\rho(\Phi_0(z')) \left.\frac{\mathrm{dNo}_t(\Phi_{-t}(z'))}{\mathrm{d}t}\right|_{t=0} = -\rho(\Phi_0(z'))\mathrm{No}_0(z')\lambda(\Phi_0(z')) \tag{49}$$

$$= -\rho(z')\lambda(z'), \tag{50}$$

using Equation (22). In summary, we have:

$$\frac{\mathrm{d}}{\mathrm{d}t} I_0(z', t)\Big|_{t=0} = \rho(z')\left\langle \nabla U(x'), v' \right\rangle - \rho(z')\lambda(z').$$

Exactly one bounce in the interval. From Equation (34), the trajectory consists in a bounce at a time T_1, occurring with density (expressed as before as a function of the final point z') $q(t_1; F^{-1}_{t,t_1}(z'))$, followed by no bounce in the interval $(T_1, t]$, an event of probability:

$$\begin{aligned} \mathrm{No}_{t-t_1}(C \circ \Phi_{t_1}(z)) &= \mathrm{No}_{t-t_1}(C \circ \Phi_{t_1} \circ F^{-1}_{t,t_1}(z')) && (51) \\ &= \mathrm{No}_{t-t_1}(\Phi_{t_1-t}(z')), && (52) \end{aligned}$$

where we used that $C^{-1} = C$. This yields:

$$I_1(z', t) = \int_0^t q(t_1; F^{-1}_{t,t_1}(z'))\rho(\Psi^{-1}_{t_1:1,t}(z'))\mathrm{No}_{t-t_1}(\Phi_{t_1-t}(z'))\mathrm{d}t_1. \quad (53)$$

To compute the derivative of the above equation at zero, we use again Leibniz's rule:

$$\frac{\mathrm{d}}{\mathrm{d}t} I_1(z', t)\Big|_{t=0} = \rho(C(z'))\lambda(C(z')).$$

Two or more bounces in the interval. For a number of bounce $n \geq 2$, we get:

$$I_n(z', t) = \int_0^t \Bigg[\underbrace{\int \cdots \int_{t_{2:n}: t_1 < t_2 \cdots < t_n < t} \rho(\Psi^{-1}_{t_{1:n},t}(z'))\widetilde{q}(t_{1:n}; t, \Psi^{-1}_{t_{1:n},t}(z'))\mathrm{d}t_{2:n}}_{\tilde{I}(t_1,t,z')} \Bigg] \mathrm{d}t_1, \quad (54)$$

and hence, using Leibniz's rule on the integral over t_1:

$$\frac{\mathrm{d}}{\mathrm{d}t} I_n(z', t)\Big|_{t=0} = \tilde{I}(0, 0, z') = 0. \quad (55)$$

Putting all terms together. Putting everything together, we obtain:

$$\frac{\mathrm{d}\mu_t(z')}{\mathrm{d}t}\Big|_{t=0} = \rho(z') \langle \nabla U(x'), v' \rangle \underbrace{-\rho(z')\lambda(z') + \rho(C(z'))\lambda(C(z'))}. \quad (56)$$

From the expression of $\lambda(\cdot)$, we can rewrite the two terms above the brace as follows:

$$\begin{aligned} & -\rho(z')\lambda(z') + \rho(C(z'))\lambda(C(z')) \\ =& -\rho(z')\lambda(z') + \rho(z')\lambda(C(z')) \\ =& -\rho(z')\max\{0, \langle \nabla U(x'), v' \rangle\} + \rho(C(z'))\max\{0, \langle \nabla U(x'), R(x')\, v' \rangle\} \\ =& -\rho(z')\max\{0, \langle \nabla U(x'), v' \rangle\} + \rho(z')\max\{0, \langle \nabla U(x'), R(x')\, v' \rangle\} \\ =& -\rho(z')\max\{0, \langle \nabla U(x'), v' \rangle\} + \rho(z')\max\{0, -\langle \nabla U(x'), v' \rangle\} \\ =& -\rho(z')\langle \nabla U(x'), v' \rangle, \end{aligned}$$

where we used that $\rho(z') = \rho(C(z'))$, $\langle \nabla U(x'), R(x')\, v' \rangle = -\langle \nabla U(x'), v' \rangle$ and $-\max\{0, f\} + \max\{0, -f\} = -f$ for any function f. Hence we have $\frac{\mathrm{d}\mu_t(z')}{\mathrm{d}t}\Big|_{t=0} = 0$, establishing that that the bouncy particle sampler $\lambda^{\mathrm{ref}} = 0$ admits ρ as invariant distribution. The invariance for $\lambda^{\mathrm{ref}} > 0$ then follows from Lemma 1 given below.

Lemma 1. *Suppose P_t is a continuous time Markov kernel and Q is a discrete time Markov kernel which are both invariant with respect to μ. Suppose we construct for $\lambda^{ref} > 0$ a Markov process $\hat{P}_t$ as follows: at the jump times of an independent Poisson process with intensity λ^{ref} we make a transition with Q and then continue according to P_t, then $\hat{P}_t$ is also μ-invariant.*

Proof. The transition kernel is given by

$$\begin{aligned} \hat{P}_t &= e^{-\lambda t}P_t + \int_0^t \mathrm{d}t_1 \lambda e^{\lambda t_1} e^{-\lambda(t-t_1)} P_{t-t_1} Q P_{t_1} \\ &\quad + \int_0^t \mathrm{d}t_1 \int_{t_1}^{t_2} \mathrm{d}t_2 \lambda^2 e^{\lambda t_1} e^{\lambda(t_2-t_1)} e^{-\lambda(t-t_2)} P_{t-t_2} Q P_{t_2-t_1} Q P_{t_1} + \ldots \end{aligned}$$

Therefore

$$\begin{aligned}\mu\hat{P}_t &= \mu\left(e^{-\lambda t}+\lambda t e^{-\lambda t}+\frac{(\lambda t)^2}{2}e^{-\lambda t}\dots\right)\\ &= \mu.\end{aligned}$$

Hence $\hat{P}_t$ is μ-invariant. □

Infinitesimal generator and irreversibility.

Next, an argument similar to the proof of π-invariance without refreshment yields the following expression for the generator of the process with refreshment:

$$\begin{aligned}&\mathbb{E}\left[h(x(t),v(t))|\,(x(0),v(0))=(x,v)\right] \qquad (57)\\ =\;& \exp\left(-\int_0^t\left(\lambda_0(s)+\lambda^{\mathrm{ref}}\right)\mathrm{d}s\right)h(x+tv,v)\\ &+\int_0^t\frac{\mathrm{d}t_1\left(\lambda_0(t_1)+\lambda^{\mathrm{ref}}\right)\cdot\lambda^{\mathrm{ref}}}{\lambda_0(t_1)+\lambda^{\mathrm{ref}}}\exp\left(-\int_0^{t_1}\left(\lambda_0(s_1)+\lambda^{\mathrm{ref}}\right)\mathrm{d}s_1\right)\int\psi(\mathrm{d}v_1)\exp\left(-\int_{t_1}^t\left(\lambda_1(s_2)+\lambda^{\mathrm{ref}}\right)\mathrm{d}s_2\right)\\ &h(x+t_1v+(t-t_1)v_1)+\\ &\int_0^t\lambda_0(t_1)\exp\left(-\int_0^{t_1}\left(\lambda_0(s_1)+\lambda^{\mathrm{ref}}\right)\mathrm{d}s_1\right)\int_{\mathbb{R}^d}\delta_{v-\frac{2\langle\nabla U,v\rangle}{\|\nabla U\|^2}v}(\mathrm{d}v_1)\exp\left(-\int_{t_1}^t\left(\lambda_1(s_2)+\lambda^{\mathrm{ref}}\right)\mathrm{d}s_2\right)\\ &h(x+t_1v+(t-t_1)v_1)\\ &+\dots\end{aligned}$$

where $\lambda_0(s)=\max\{0,\langle\nabla U(x+sv),v\rangle\}$ and $\lambda_1(s)=\max\{0,\langle\nabla U(x+t_1v+(s-t_1)v_1),v_1\rangle\}$ are the intensity of bounces. Here we used that a Poisson process with rate $\lambda_0(t_1)+\lambda^{\mathrm{ref}}$ have bounce with probability $\frac{\lambda_0(t_1)}{\lambda_0(t_1)+\lambda^{\mathrm{ref}}}$ and refreshment with probability $\frac{\lambda^{\mathrm{ref}}}{\lambda_0(t_1)+\lambda^{\mathrm{ref}}}$. From this we can conclude that for $(x(0),v(0))=(x,v)$

$$\begin{aligned}\mathcal{L}h &= \lim_{t\to0}\frac{\mathbb{E}\left[h(x(t),v(t))|\,(x(0),v(0))=(x,v)\right]-h(x,v)}{t}=-\left(\max\{0,\langle\nabla U(x),v\rangle\}\right)h(x,v)+\langle\nabla_x h,v\rangle\\ &\quad+\lambda^{\mathrm{ref}}\int\psi(\mathrm{d}v')\left(h(x,v')-h(x,v)\right)+\max\{0,\langle\nabla U(x),v\rangle\}\,h(x,R(x)v).\end{aligned}$$

A simple calculation shows that

$$\int\mathrm{d}z\rho\left((\mathcal{L}h)\,g-(\mathcal{L}g)\,h\right)=\int\mathrm{d}z\rho(z)\left(h(x,v-\frac{2\langle\nabla U,v\rangle}{\|\nabla U\|^2}v)g-hg\right)\langle\nabla U(x),v\rangle.$$

Generally, the above expression is not equal to zero, hence the stochastic process is not reversible.

A.2 Ergodicity

We denote by ${P^{\mathrm{ref}}}_t(z,\mathrm{d}z')$ the transition kernel of the continuous-time Markov process corresponding to the potential $U=0$, that is we consider the process without bounces based purely on refreshments at constant rate. We will compare ${P^{\mathrm{ref}}}_t$ to P_t in order to show that ρ is the unique invariant distribution of P_t. We fix $t_0>0$ and consider the embedded discrete time chain ${P^{\mathrm{ref}}}_{nt_0}$. In the following, $B_R(x)$denotes the Euclidean ball of radius R centered at x. Our proof relies on the following preliminary lemma.

Lemma 2. *${P^{ref}}_{t_0}$ satisfies that*

1. *for all $t_0>0$, $z=(x,v)\in\mathbb{R}^d\times\mathbb{R}^d$ and open set $W\subset\mathbb{R}^d\times\mathbb{R}^d$ there exists $n\in\{1,2,3,\dots\}$ such that ${P^{ref}}_{nt_0}(z,W)>0$*

2. *For all* $\epsilon > 0$ *there exists* $\delta > 0$ *and* $t_0 > 0$ *such that for any* $x_0 \in \mathbb{R}^d$ *and any* $(x, v) \in B_\epsilon(x_0) \times B_1(0)$,

$$P^{ref}{}_{t_0}\left((x,v), \mathrm{d}\left(x', v'\right)\right) \geq \delta \mathbb{1}_{B_\epsilon(x_0)\times B_1(0)}\left(\mathrm{d}x, \mathrm{d}v\right). \tag{58}$$

Proof of Lemma 2. We only prove 2) as 1) can be established in a similar way. Let h be a positive test function. The aim is to show that for $z \in B_\epsilon(\hat{x}) \times B_1(0)$

$$P^{\mathrm{ref}}{}_t h(z) \geq \delta \int_{B_\epsilon(x_0)\times B_1(0)} h(z')\mathrm{d}z'. \tag{59}$$

We pick $t = 4\epsilon$ and obtain a non explicit lower bound for δ. The transition kernel $P^{\mathrm{ref}}{}_t\left(z, \mathrm{d}z'\right)$ is bounded from below by considering the transition due to exactly two refreshments, that is

$$\begin{aligned} P^{\mathrm{ref}}{}_t g(z) \geq & \int_{\left\{0<t_1,t_2<\frac{\epsilon}{2}, t_1+t_2<t\right\}} \mathrm{d}t_1\mathrm{d}t_2 \int_{\mathbb{R}^d} \psi(\mathrm{d}v_1) \int_{\mathbb{R}^d} \psi(\mathrm{d}v_2) \\ & {\lambda^{\mathrm{ref}}}^2 \exp\left(-\lambda^{\mathrm{ref}} t\right) h\left(x + t_0 v + t_1 v_1 + (t - t_1 - t_0)v_2, v_2\right). \end{aligned} \tag{60}$$

To obtain a lower bound as in Equation (58), we use the co-area formula.

For C^1 Riemannian manifolds M and N of dimension m and n and a differentiable map F and h a test function, we have

$$\begin{aligned} \int_M h(w) JF(w) \mathrm{d}\mathcal{H}_m(w) &= \int_N \mathrm{d}\mathcal{H}_n(u) \int_{F^{-1}(u)} \mathrm{d}\mathcal{H}_{m-n}(x) h(x), \\ \int_M \frac{h(F(w))}{JF(w)} JF(w) \mathrm{d}\mathcal{H}_m(w) &= \int_N \mathrm{d}\mathcal{H}_n(u) h(u) \int_{F^{-1}(u)} \mathrm{d}\mathcal{H}_{m-n}(x) \frac{1}{JF(x)}, \end{aligned}$$

see e.g. Section 3.2 of [1]. Here $\mathcal{H}_m$, $\mathcal{H}_n$ and $\mathcal{H}_{m-n}$ denote the volume measures associated with Riemannian metric on M, N and $F^{-1}(u)$ (with the induced metric of M). In the above equations, JF is a generalization of the determinant of the Jacobian $JF := \sqrt{\det g_t(\nabla f_i, \nabla f_j)}$ where g is the corresponding Riemannian metric, see [1]. Here $JF = \sqrt{\det DF'\ DF}$. Applying the statement to Equation (60) with $M = \left\{0 < t_1, t_2 < \frac{\epsilon}{2}, t_1 + t_2 < t\right\} \times \mathbb{R}^d \times \mathbb{R}^d$ and $N = \mathbb{R}^d \times \mathbb{R}^d$,

$$\begin{aligned} P^{\mathrm{ref}}{}_t g(z) \geq {\lambda^{\mathrm{ref}}}^2 \exp\left(-\lambda^{\mathrm{ref}} t\right) \int h\left(z_2\right) \mathrm{d}z_2 \\ \int_{F^{-1}_{t,x_0,v_0}(z_2)\cap M} \mathrm{d}\mathcal{H}(t_1, t_2, v_1, v_2) \psi(v_1) \int_{\mathbb{R}^d} \psi(v_2) \frac{1}{JF_{t,x_0,v_0}(t_1, t_2, v_1, v_2)}. \end{aligned}$$

with

$$\begin{aligned} F_{t,x_0,v_0}(t_0, v_1, t_1, v_2) :&= \left(x_0 + t_0 v_0 + t_1 v_1 + (t - t_1 - t_0)v_2, v_2\right). \\ DF_{t,x_0,v_0} &= \begin{pmatrix} v_0 & t_1 I & v_1 - v_2 & (t - t_0 - t_1) I \\ 0 & 0 & 0 & I \end{pmatrix} \\ JF_{t,x_0,v_0} &= \sqrt{\det DF_{t,x_0,v_0} DF^t_{t,x_0,v_0}}. \end{aligned}$$

Notice that $F^{-1}_{t,x_0,v_0}(z_2) \neq \emptyset$ because $\{t_0, v_1, t_1, v_2\} \in F^{-1}_{t,x_0,v_0}(z_2)$ for $0 < t_0 < \frac{\epsilon}{2}$, $t_1 < t - t_0 - \frac{\epsilon}{2}$ and $v_1 = \frac{(x_2 - (t - t_0 - t_1)v_2) - (x_0 + t_0 v_0)}{t_1}$.

The statement can now be verified by showing that

$$\inf_{z, z_2} \int_{F_z^{-1}(z_2)\cap M} \mathrm{d}\mathcal{H}(t_1, t_2, v_1, v_2) \frac{1}{JF(t_1, t_2, v_1, v_2)} > 0,$$

which follows from the fact that DF_{t,x_0,v_0} has full rank on M. □

Proof of Theorem 1. Suppose BPS is not ergodic, then it follows from standard results in ergodic theory that there are two measures μ_1 and μ_2 such that $\mu_1 \perp \mu_2$ and $\mu_i P_{t_0} = \mu_i$, see e.g. [10]. Thus there is a set A such that

$$\mu_1(A) = \mu_2(A^c) = 0 \tag{61}$$

Because of the first part of Lemma 2 and Lemma 2.2 of [10] the support of the μ_i is $\mathbb{R}^d \times \mathbb{R}^d$. Thus

$$\infty > \mu_i \left(B_\epsilon(0) \times B_1(0) \right) > \tilde{\delta} > 0$$

and we have by invariance

$$\begin{aligned}
\mu_i(\mathrm{d}\tilde{z}) &= \int \mu_i \left(\mathrm{d}z \right) P_{t_0}(z, \mathrm{d}\tilde{z}) \\
&\geq \int_{B_\epsilon(0) \times B_1(0)} \mu_i \left(\mathrm{d}z \right) \exp\left(-t_0 \sup_{w \in B_{2\epsilon}(0)} \|\nabla U(w)\| \right) P^{\mathrm{ref}}{}_{t_0}(z, \mathrm{d}\tilde{z}) \\
&\geq \mathbb{1}_{B_\epsilon(0) \times B_1(0)} \delta \exp\left(- \sup_{w \in B_{2\epsilon}(0)} \|\nabla U(w)\| \right) \mathrm{d}z,
\end{aligned} \tag{62}$$

where the last step follows by Lemma 2 for some $\delta > 0$. That contradicts the fact that $\mu_1 \perp \mu_2$.

The law of large numbers then follows by Birkhoff's pointwise ergodic theorem; see e.g. [7, Theorem 2.30, Section 2.6.4]. □

A.3 Local Sampler

In the following we adapt the computations of Appendix A.1 to show invariance for the local sampler. Similar to the basic BPS it is sufficient to compute I_0 and I_1 given in (43) and (53) because the higher order terms have a vanishing derivative. The local sampler can be realized by considering the global events Q, see Algorithm 4, as coming from a Poisson process with intensity $\lambda^{\mathrm{g}} = \sum_{f \in F} \lambda_f$ a bounce event at time t is then associated with factor f with probability $\frac{\lambda_f(z(t))}{\sum_{\tilde{f} \in F} \lambda_{\tilde{f}}(z(t))}$.

We obtain for the term I_0 corresponding to no bounce happening

$$\rho\left(\Phi_{-t}\left(z' \right) \right) \mathrm{No}_t\left(\Phi_{-t}(z') \right),$$

where $\mathrm{No}^{\mathrm{g}}_t\left(\Phi_{-t}\left(\tilde{z} \right) \right)$ where is the probability that no bounce happens on $[0, t]$ when started at $\tilde{z}$, that is

$$\mathrm{No}_t\left(\Phi_{-t}\left(\tilde{z} \right) \right) = \exp\left(- \int_0^t \Phi_s\left(\tilde{z} \right) \mathrm{d}s \right).$$

We now evaluate the time derivative at zero of $I_0\left(z', t \right)$

$$\frac{\mathrm{d}}{\mathrm{d}t} I_0\left(z', t \right) \mid_{t=0} = \rho\left(z' \right) \langle \nabla U(x), v \rangle - \rho(z') \lambda^{\mathrm{g}}\left(z' \right).$$

The term $I_1\left(z, t \right)$ is obtained by performing a change of coordinates

$$\begin{aligned}
&\mathbb{E}\left[\mathbb{E}\left[h(z\left(t \right)) 1_{\#\mathrm{Col}_t(T)=1} \middle| z\left(0 \right) = z \right] \right] \\
&= \int_0^t \mathrm{d}t_1 \int \mathrm{d}z \rho(z) \left(\sum_{f' \in F} \left(\lambda_{f'} \circ \Phi_{t_1} \right)(z) \right) \mathrm{No}_{t-t_1}(C_f \circ \Phi_{t_1}(z)) \sum_{f \in F} \frac{\left(\lambda_f \circ \Phi_{t_1} \right)(z)}{\sum_{f' \in F} \left(\lambda_{f'} \circ \Phi_{t_1} \right)(z)} \\
&\quad \cdot \left(h \circ \Phi_{t-t_1} \circ C_f \circ \Phi_{t_1} \right)(z) \\
&= \sum_{f \in F} \int_0^t \mathrm{d}t_1 \int \mathrm{d}z \rho\left(z \right) \mathrm{No}_{t-t_1}(C_f \circ \Phi_{t_1}(z)) \lambda_f \left(h \circ \Phi_{t-t_1} \circ C_f \circ \Phi_{t_1} \right)(z),
\end{aligned}$$

where

$$C_f\left((x,v)\right) = (x, R_f(x)v)\,.$$

Performing a change of coordinates for each of the terms and the sum, we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t} I_1\left(z',t\right) \mid_{t=0} = \sum_{f\in F} \rho(C_f(z'))\lambda(C_f(z')).$$

Thus, we obtain overall

$$\frac{\mathrm{d}}{\mathrm{d}t}\mu_t(z') \mid_{t=0} = \rho\left(z'\right) \sum_{f\in F} \langle \nabla U_f(x), v\rangle - \rho(z') \sum_{f\in F} \max\left\{0, \langle \nabla U_f, v\rangle\right\} + \sum_{f\in F} \rho(C_f(z')) \sum_{f\in F} \max\left\{0, -\langle \nabla U_f, v\rangle\right\}$$

Inserting ρ from Equation (12), we obtain $\frac{\mathrm{d}}{\mathrm{d}t}\mu_t \mid_{t=0} = 0$.

B Bayesian logistic regression for large datasets

B.1 Bounds on the intensity

We derive here a datapoint-specific upper bound $\bar{\chi}^{[r]}$ to $\chi^{[r]}(t)$. First, we need to compute the gradient for one datapoint:

$$\nabla U^{[r]}(x) = \iota_r(\mathrm{logistic}\left\langle \iota_r, x\right\rangle - c_r),$$

where:

$$\mathrm{logistic}(a) = \frac{e^a}{1+e^a}.$$

We then consider two sub-cases depending on $c_r = 0$ or $c_r = 1$. Suppose first $c_r = 0$, and let $x(t) = x + tv$

$$\begin{aligned}
\chi^{[r]}(t) &= \max\{0, \langle \nabla U_r(x(t)), v\rangle\} \\
&= \max\left\{0, \sum_{k=1}^{d} \iota_{r,k} v_k \mathrm{logistic}\left\langle \iota_r, x(t)\right\rangle\right\} \\
&\leq \sum_{k=1}^{d} \mathbf{1}[v_k > 0] \iota_{r,k} v_k \mathrm{logistic}\left\langle \iota_r, x(t)\right\rangle \\
&\leq \underbrace{\sum_{k=1}^{d} \mathbf{1}[v_k > 0] \iota_{r,k} v_k}_{\bar{\chi}^{[r]}}
\end{aligned}$$

The bound makes sense intuitively: it the true label is $c_r = 0$, then to correctly label this datapoint, for the dimensions where $\iota_{r,k}$ is large, we need small values of x_k. This means we want the bouncing rate to increase with the velocity for such dimension.

Similarly, for $c_r = 1$:

$$\begin{aligned}
\chi^{[r]}(t) &= \max\left\{0, \sum_{k=1}^{d} \iota_{r,k} v_k (\text{logistic}\left\langle \iota_r, x(t)\right\rangle - 1)\right\} \\
&= \max\left\{0, \sum_{k=1}^{d} \iota_{r,k} (-v_k)(1 - \text{logistic}\left\langle \iota_r, x(t)\right\rangle)\right\} \\
&\leq \sum_{k=1}^{d} \mathbf{1}[v_k < 0] \iota_{r,k} (-v_k)(1 - \text{logistic}\left\langle \iota_r, x(t)\right\rangle) \\
&= \sum_{k=1}^{d} \mathbf{1}[v_k < 0] \iota_{r,k} |v_k| (1 - \text{logistic}\left\langle \iota_r, x(t)\right\rangle) \\
&\leq \underbrace{\sum_{k=1}^{d} \mathbf{1}[v_k < 0] \iota_{r,k} |v_k|}_{\bar{\chi}^{[r]}}
\end{aligned}$$

Combining these we obtain:

$$\bar{\chi}^{[r]} = \sum_{k=1}^{d} \mathbf{1}[v_k (-1)^{c_r} \geq 0] \iota_{r,k} |v_k|.$$

When implementing Algorithm 6, we need to bound $\sum_{r=1}^{R} \chi^{[r]}(t)$. We have

$$\begin{aligned}
\bar{\chi} &= \sum_{r=1}^{R} \bar{\chi}^{[r]} \geq \sum_{r=1}^{R} \chi^{[r]}(t) \\
&= \sum_{r=1}^{R} \sum_{k=1}^{d} \mathbf{1}[v_k (-1)^{c_r} \geq 0] \iota_{r,k} \ |v_k| \\
&= \sum_{k=1}^{d} |v_k| \sum_{r=1}^{R} \mathbf{1}[v_k (-1)^{c_r} \geq 0] \iota_{r,k} \\
&= \sum_{k=1}^{d} |v_k| \ \iota_k^{\mathbf{1}[v_k < 0]}.
\end{aligned}$$

Remark 2. Notice that above is uniformly bounded between events with the bound only depending on the magnitude of v. Alternative by choosing restricted refreshment a bound for all time can be derived which allows the implementation of the subsampling procedure outlined in Remark 1.

B.2 Sampling the thinned factor

We show here how to implement Step 5(c)ii of Algorithm 6 without enumerating over the R datapoints. We begin by introducing some required pre-computed data-structures. The pre-computation is executed only once at the beginning of the algorithm, so its running time, $O(R \log R)$ is considered negligible (the number of bouncing events is assumed to be greater than R). For each dimensionality k and class label c, consider the categorical distribution with the following probability mass function over the datapoints:

$$\mu_k^c(r) = \frac{\iota_{r,k} \mathbf{1}[c_r = c]}{\iota_k^c}.$$

This is just the distribution over the datapoints that have the given label, weighted by the covariate k. An alias sampling data-structure [6, Section 3.4] is computed for each k and c. This pre-computation takes total time $O(R \log R)$. This allows subsequently to sample in time $O(1)$ from the distributions μ_k^c.

We now show how this pre-computation is used to to implement Step 5(c)ii of Algorithm 6. We denote the probability mass function we want to sample from by:

$$q(r) = \frac{\bar{\chi}^{[r]}}{\bar{\chi}}.$$

To do this sampling efficiently, we construct an artificial joint distribution over both datapoints and covariate dimension indices:

$$q(r,k) = \frac{\mathbf{1}[v_k(-1)^{c_r} \geq 0]\iota_{r,k}\ |v_k|}{\bar{\chi}}.$$

We denote by $q_{\mathrm{k}}(k)$, respectively $q_{\mathrm{r|k}}(r|k)$, the associated marginal, respectively conditional distribution. By construction, we have

$$\sum_{k=1}^{d} q(r,k) = q(r).$$

It is therefore enough to sample (r,k) and to return r. To do so, we first sample (a) $k \sim q_{\mathrm{k}}(\cdot)$ and then (b) sample $r|k \sim q_{\mathrm{r|k}}(\cdot|k)$.

For (a), we have:

$$q_{\mathrm{k}}(k) = \frac{|v_k|\ \iota_k^{\mathbf{1}[v_k<0]}}{\bar{\chi}},$$

so this sampling step again does not require looping over the datapoints thanks the pre-computations described earlier.

For (b), we have:

$$q_{\mathrm{r|k}}(r|k) = \mu_k^{\mathbf{1}[v_k<0]}(r),$$

and therefore this sampling step can be computed in $O(1)$ thanks to the pre-computed alias sampling data-structure.

B.3 Algorithm description

Algorithm 7 contains a detailed implementation of the local BPS with thinning for the logistic regression example (Example 4.6).

Algorithm 7 Local BPS algorithm for Logistic Regression with Large Datasets

1. Precompute the alias tables $\mu_k^c(r)$ for $k = 1, \dots, d$, $c \in \{0, 1\}$ in order to sample from $r \sim q_{\mathrm{r|k}}(\cdot|k)$.
2. Initialize the state and velocity $\left(x^{(0)}, v^{(0)}\right)$ arbitrarily on $\mathbb{R}^d \times \mathbb{R}^d$.
3. Initialize the global clock $T \leftarrow 0$.
4. Initialize $\bar{T} \leftarrow \triangle$. (time until which local upper bounds are valid)
5. Compute local-in-time upper bound on the prior factor as $\bar{\chi}_{\mathrm{prior}} = \sigma^{-2} \max\left(0, \left\langle x^{(0)} + v^{(0)}\Delta, v^{(0)} \right\rangle\right)$ (notice the rate associated with the prior is monotonically increasing).
6. While more events $i = 1, 2, \dots$ requested do
   (a) Compute the local-in-time upper bound on the data factors in $\mathcal{O}(d)$
   $$\bar{\chi} = \sum_{k=1}^{d} |v_k^{(i-1)}| \, \iota_k^{\mathbf{1}[v_k^{(i-1)} < 0]}.$$
   (b) Sample $\tau \sim \mathrm{Exp}\left(\bar{\chi}_{\mathrm{prior}} + \bar{\chi} + \lambda^{\mathrm{ref}}\right)$.
   (c) If $\left(T + \tau > \bar{T}\right)$ then
      i. $x^{(i)} \leftarrow x^{(i-1)} + v^{(i-1)}(\bar{T} - T)$.
      ii. $v^{(i)} \leftarrow v^{(i-1)}$.
      iii. Compute local-in-time upper bound $\bar{\chi}_{\mathrm{prior}} = \sigma^{-2} \max\left(0, \left\langle x^{(i)} + v^{(i)}\Delta, v^{(i)} \right\rangle\right)$.
      iv. Set $T \leftarrow \bar{T}$, $\bar{T} \leftarrow \bar{T} + \triangle$.
   (d) Else
      i. $x^{(i)} \leftarrow x^{(i-1)} + v^{(i-1)}\tau$.
      ii. Sample j from Discrete$\left(\frac{\bar{\chi}}{\bar{\chi}_{\mathrm{prior}} + \bar{\chi} + \lambda^{\mathrm{ref}}}, \frac{\lambda^{\mathrm{ref}}}{\bar{\chi}_{\mathrm{prior}} + \bar{\chi} + \lambda^{\mathrm{ref}}}, \frac{\bar{\chi}_{\mathrm{prior}}}{\bar{\chi}_{\mathrm{prior}} + \bar{\chi} + \lambda^{\mathrm{ref}}}\right)$.
      iii. If $j = 1$
         A. Sample k according to $q_{\mathrm{k}}(k) = |v_k^{(i-1)}| \, \iota_k^{\mathbf{1}[v_k^{(i-1)} < 0]} / \bar{\chi}$.
         B. Sample $r \sim q_{\mathrm{r|k}}(\cdot|k)$ using the precomputed alias table.
         C. If $V < \frac{\max\left(0, \left\langle \nabla U^{[r]}(x^{(i)}), v^{(i-1)} \right\rangle\right)}{\bar{\chi}^{[r]}}$ where $V \sim \mathcal{U}(0, 1)$.
         $v^{(i)} \leftarrow R_r(x^{(i)}) v^{(i-1)}$ where R_r is the collision operator associated with the r-th data item .
         D. Else $v^{(i)} \leftarrow v^{(i-1)}$.
      iv. If $j = 2$
         A. $v^{(i)} \sim \mathcal{N}(0_d, I_d)$.
      v. If $j = 3$
         A. If $V < \frac{\sigma^{-2} \max\left(0, \left\langle x^{(i)}, v^{(i-1)} \right\rangle\right)}{\chi_{\mathrm{prior}}}$ where $V \sim \mathcal{U}(0, 1)$.
         $v^{(i)} \leftarrow R_{\mathrm{prior}}(x^{(i)}) v^{(i-1)}$ where R_{prior} is collision operator associated with prior.
         B. Else $v^{(i)} \leftarrow v^{(i-1)}$.
      vi. Compute local-in-time upper bound $\bar{\chi}_{\mathrm{prior}} = \sigma^{-2} \max\left(0, \left\langle x^{(i)} + v^{(i)}\left(\bar{T} - T - \tau\right), v^{(i)} \right\rangle\right)$.
      vii. Set $T \leftarrow T + \tau$

C Calculations of Section 4.1 and proof of Proposition 2

C.1 Calculations of the norms at bounce times in the isotropic normal case

From Equation (6), we have

$$\begin{aligned}\left\langle x^{(i)}, v^{(i)}\right\rangle &= \left\langle x^{(i-1)}+\tau_i v^{(i-1)}, v^{(i-1)} - \frac{2\left\langle x^{(i-1)}+\tau_i v^{(i-1)}, v^{(i-1)}\right\rangle}{\left\|x^{(i-1)}+\tau_i v^{(i-1)}\right\|^2}\left(x^{(i-1)}+\tau_i v^{(i-1)}\right)\right\rangle \\ &= \left\langle x^{(i-1)}, v^{(i-1)}\right\rangle + \tau_i - 2\left\langle x^{(i-1)}+\tau_i v^{(i-1)}, v^{(i-1)}\right\rangle = -\left\langle x^{(i-1)}, v^{(i-1)}\right\rangle - \tau_i \\ &= \begin{cases} -\sqrt{-\log V_i} & \text{if } \left\langle x^{(i-1)}, v^{(i-1)}\right\rangle \leq 0 \\ -\sqrt{\left\langle x^{(i-1)}, v^{(i-1)}\right\rangle^2 - \log V_i} & \text{otherwise} \end{cases},\end{aligned}$$

and we can obtain easily the expression for $\left\|x^{(i)}\right\|^2$ shown in 4.1.

C.2 Proof of Proposition 2

The dynamics of the BPS can be lumped into a two dimensional Markov process involving only the radius $r_t = \|x(t)\|$ and $m_t = \langle x(t), v(t)\rangle / \|x(t)\|$ for any dimensionality $d \geq 2$. The variable m can be interpreted (via $\arccos(m)$) as the angle between the particle position and velocity. Because of the strong Markov property we take $\tau_1 = 0$ without loss of generality and let t be some time between the current event and the next, yielding:

$$\begin{aligned} r_t &= \sqrt{\langle x(0)+v(0)\cdot t, x(0)+v(0)\cdot t\rangle} = \sqrt{r_0^2 + 2t_0 m(0) r_0 + t^2} \\ m_t &= \frac{\langle x(0)+v(0)\cdot t, v(0)\rangle}{\|x(0)+v(0)\cdot t\|} = \frac{r_0 m(0)+t}{r_t} \end{aligned} \tag{63}$$

If there is a bounce at time t, then

$$\begin{aligned} R_t' &= R_t \\ M_t' &= -M_t. \end{aligned} \tag{64}$$

The bounce happens with intensity $\lambda(x+tv, v) = \max(0, \langle x, v\rangle + t)$. These processes can also be written as an stochastic differential equation (SDE) driven by a jump process whose intensity is coupled to its position. This is achieved by writing the deterministic dynamics given in (63) between events as the following ordinary differential equation (ODE):

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} r_t &= \frac{2m_t r_t}{2r_t} = m_t \\ \frac{\mathrm{d}}{\mathrm{d}t} m_t &= \frac{r_t - (r_t m_t) m_t}{(r_t)^2} = \frac{1-(m_t)^2}{r_t}. \end{aligned}$$

Taking the bounces into account turns this ODE into an SDE with

$$\begin{aligned} \mathrm{d}R_t &= \frac{2M_t R_t}{2R_t} = M_t \\ \mathrm{d}M_t &= \frac{1-M_t^2}{R_t}\,\mathrm{d}t - 2M_t\,\mathrm{d}N_t \end{aligned} \tag{65}$$

where N_t is the counting process associated with Π_t the inhomogeneous Poisson process with intensity $\max(0, R_t \cdot M_t)$.

Now consider the push forward measure of $\mathcal{N}\left(0, \frac{1}{2} I_k\right) \otimes \mathcal{U}(\mathcal{S}^{k-1})$ under the map $(x, v) \mapsto (\|x\|, \langle x, v\rangle / \|x\|)$ where $\mathcal{U}(S^{k-1})$ is the uniform distribution on $\mathcal{S}^{k-1}$. This yields the collection of measures with densities $f_k(r, m)$. It is easy to check that $f_k(r, m)$ is invariant for (65) for all $k \geq 2$.

C.3 Markov chain on the events

To facilitate this investigation, we first use Equations (63) and (64) to obtain a discrete time Markov chain on the events. An algorithm to perform this simulation is shown in Algorithm 8.

Algorithm 8 Event Chain for isotropic Gaussian target distributions.

---

1. Set $R_0 \leftarrow 0$ and $M_0 \leftarrow 0$
2. For $i = 1, \dots, n$
   (a) Simulate $b_i \sim \text{Exp}(1), \tau_{\text{ref}i} \sim \text{Exp}(1), Y_i \sim \text{Beta}(\frac{d-1}{2}, \frac{d-1}{2})$
   (b) Set $\tau_{ci} \leftarrow \begin{cases} -M_{i-1} + \sqrt{b_i} & \text{for } M_{i-1} \leq 0 \\ -M_{i-1} + \sqrt{M_{i-1}^2 + b_i} & \text{otherwise} \end{cases}$
   (c) If $\tau_{ci} \leq \tau_{\text{ref}i}$
   - $R_i \leftarrow \sqrt{R_{i-1}^2 + 2\tau_{ci} M_{i-1} R_{i-1} + {\tau_c}^2}$
   - $M_i \leftarrow \begin{cases} -\frac{\sqrt{b_i}}{R_i} & M_{i-1} < 0 \\ -\frac{\sqrt{b_i + M_{i-1}^2 R_{i-1}^2}}{R_i} \end{cases}$

   (d) Else
   - $R_i \leftarrow \sqrt{R_{i-1}^2 + 2\tau_{\text{ref}i} M_{i-1} R_{i-1} + \tau_{\text{ref}i}}$
   - $M_i \leftarrow 2Y_i - 1$